Representation theory

E. Kowalski

ETH Zürich – Spring Semester 2011
Version of September 23, 2011
kowalski@math.ethz.ch

Contents

Chapter 1. Introduction and motivation 1
  1.1. Presentation 2
  1.2. Four motivating statements 3
  1.3. Prerequisites and notation 5

Chapter 2. The language of representation theory 8
  2.1. Basic language 8
  2.2. Formalism: changing the space 13
  2.3. Formalism: changing the group 26
  2.4. Formalism: changing the field 42
  2.5. Matrix representations 44
  2.6. Examples 45
  2.7. Some general results 52
  2.8. Conclusion 81

Chapter 3. Variants 83
  3.1. Representations of algebras 83
  3.2. Topological groups 86
  3.3. Unitary representations 91

Chapter 4. Linear representations of finite groups 99
  4.1. Maschke's Theorem 99
  4.2. Applications of Maschke's Theorem 102
  4.3. Decomposition of representations 106
  4.4. Harmonic analysis on finite groups 122
  4.5. Finite abelian groups 130
  4.6. The character table 136
  4.7. Applications 156
  4.8. Further topics 173

Chapter 5. Abstract representation theory of compact groups 178
  5.1. An example: the circle group 178
  5.2. Haar measure and the regular representation of a compact group 180
  5.3. The analogue of the group algebra 191
  5.4. The Peter-Weyl theorem 196
  5.5. Characters and matrix coefficients for compact groups 203
  5.6. Some first examples 207

Chapter 6. Applications of representations of compact groups 214
  6.1. Compact Lie groups are matrix groups 214
  6.2. The Frobenius-Schur indicator 217
  6.3. The Larsen alternative 223
  6.4. The Hydrogen atom 230

Chapter 7. Other groups: a few examples 237
7.1. Algebraic groups 237
7.2. Locally-compact abelian groups 242
7.3. A non-abelian example: $\mathrm{SL}_2(\mathbf{R})$ 242

Appendix A. Some useful facts 243
A.1. Algebraic integers 243
A.2. The spectral theorem 246
A.3. The Stone-Weierstrass theorem 251

Bibliography 252

CHAPTER 1

Introduction and motivation

These notes are intended to provide a basic introduction to (some of) the fundamental ideas and results of *representation theory* of groups. In this first preliminary chapter, we start with some motivating remarks and provide a general overview of the rest of the text; we also include some notes on the prerequisite knowledge – which are not uniform for all parts of the notes – and discuss the basic notation that we use.

In writing this text, the objective has never been to give the shortest or slickest proof. To the extent that the author's knowledge makes this possible, the goal is rather to explain the ideas and the mechanism of thought that can lead to an understanding of "why" something is true, and not simply to the quickest line-by-line check that it holds.

The point of view is that representation theory is a fundamental theory, both for its own sake and as a tool in many other fields of mathematics; the more one knows, understands and breathes representation theory, the better. This style (or its most ideal potential form) is perhaps best summarized by P. Sarnak's advice in the Princeton Companion to Mathematics [**17**, p. 1008]:

> One of the troubles with recent accounts of certain topics is that they can become too slick. As each new author finds cleverer proofs or treatments of a theory, the treatment evolves toward the one that contains the "shortest proofs." Unfortunately, these are often in a form that causes the new student to ponder, "How did anyone think of this?" By going back to the original sources one can usually see the subject evolving naturally and understand how it has reached its modern form. (There will remain those unexpected and brilliant steps at which one can only marvel at the genius of the inventor, but there are far fewer of these than you might think.) As an example, I usually recommend reading Weyl's original papers on the representation theory of compact Lie groups and the derivation of his character formula, alongside one of the many modern treatments.

So the text sometimes gives two proofs of the same result, even in cases where the arguments are fairly closely related; one may be easy to motivate ("how would one try to prove such a thing?"), while the other may recover the result by a slicker exploitation of the formalism of representation theory. To give an example, we first consider Burnside's irreducibility criterion, and its developments, using an argument roughly similar to the original one, before showing how Frobenius reciprocity leads to a quicker line of reasoning (see Sections 2.7.3 and 2.7.4).

Acknowledgments. The notes were prepared in parallel with the course "Representation Theory" that I taught at ETH Zürich during the Spring Semester 2011. Thanks are obviously due to all the students who attended the course for their remarks and interest, in particular M. Lüthy, M Rüst, I. Schwabacher, M. Scheuss, and M. Tornier, and to the assistants in charge of the exercise sessions, in particular J. Ditchen who coordinated those. Thanks also "Anonymous Rex" for a comment on a blog post, to U. Schapira for

his comments and questions during the class, and to A. Venkatesh for showing me his own notes for a (more advanced) representation theory class, from which I derived much insight.

1.1. Presentation

A (linear) representation of a group G is, to begin with, simply a *homomorphism*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

where E is a vector space over some field k and $\mathrm{GL}(E)$ is the group of invertible k-linear maps on E. Thus one can guess that this should be a useful notion by noting how it involves the simplest and most ubiquitous algebraic structure, that of a group, with the powerful and flexible tools of linear algebra. Or, in other words, such a map attempts to "represent" the elements of G as symmetries of the vector space E (note that ϱ might fail to be injective, so that G is not mapped to an isomorphic group).

But even a first guess would probably not lead to imagine how widespread and influential the concepts of representation theory turn out to be in current mathematics. Few fields of pure mathematics, or of mathematical physics (or chemistry), do not make use of these ideas, and many depend on representations in an essential way. We will try to illustrate this wide influence with examples, taken in particular from number theory and from basic quantum mechanics; already in Section 1.2 below we state four results, where representation theory does not appear in the statements although it is a fundamental tool in the proofs. Moreover, it should be said that representation theory is now a field of mathematics in its own right, which can be pursued without having immediate applications in mind; it does not require external influences to expand with new questions, results and concepts – but we will barely scratch such aspects, by lack of knowledge.

The next chapter starts by presenting the fundamental vocabulary that is the foundation of representation theory, and by illustrating them with examples. It also contains a number of short remarks concerning variants of the definition of a representation ϱ: restrictions can be imposed on the group G, on the type of fields or vector spaces E allowed, or additional regularity assumptions may be imposed on ϱ when this makes sense; or even the target groups $\mathrm{GL}(E)$ may be replaced by suitable subgroups, or even by other classes of basic groups, such as symmetric groups. Many of these variants are important topics in their own right, but most of them won't reappear in the rest of these notes – lack of space, if not as well of expertise, is to blame.

Chapter 4 is an introduction to the simplest case of representation theory: the linear representations of finite groups in finite-dimensional complex vector spaces. This is also historically the first case that was studied in depth by Dirichlet (for finite abelian groups), then Frobenius, Schur, Burnside, and many others. It is a beautiful theory, and has many important applications. It also serves as "blueprint" to many generalizations, as we will see in the rest of the notes: various facts, which are extremely elementary for finite groups, remain valid, when properly framed, for important classes of infinite groups.

Among these, the compact topological groups are undoubtedly those closest to finite groups, and we consider them first abstractly in the following chapter. Then another chapter presents some concrete examples of compact Lie groups (compact matrix groups, such as unitary groups $\mathrm{U}(n, \mathbf{C})$), with some applications – the most important being probably the way representation theory explains a lot about the way the most basic atom, Hydrogen, behaves in the real world...

The final chapter is only an introduction to the next class of infinite groups, the locally compact groups, but non-compact, groups, through the fundamental examples of the abelian group $\mathbf{R}$ and the group $\mathrm{SL}_2(\mathbf{R})$ of two-by-two real matrices with determinant 1. We use this example primarily to illustrate some of the striking new phenomena that arise when compactness is lost. Here, although we can not give complete proofs, we also have an important concrete application in mind: the use of a famous theorem of Selberg to construct so-called *expander graphs.*

In an Appendix, we have gathered statements and sketches of proofs for certain facts, especially the Spectral Theorem for bounded or unbounded self-adjoint linear operators, which are needed for rigorous treatments of compact and locally compact groups. We have also included a short section illustrating concrete computations of representation theory, using such software packages as Magma or GAP.

Throughout, we also present some examples by means of exercises. These are usually not particularly difficult, but we hope they will help the reader to get acquainted with the way of thinking that representation theory often suggests for certain problems...

1.2. Four motivating statements

Below are four results, taken in very different fields, which we will prove later (or sometimes only sketch when very different ideas are also needed). The statements do not mention representation theory, in fact two of them do not even mention groups explicitly. Yet they are proved using these tools, and they serve as striking illustrations of what can be done using representation theory.

Example 1.2.1 (Primes in arithmetic progressions). Historically, the first triumph of representation theory is the proof, by Dirichlet, of the existence of infinitely many primes in arithmetic progressions, whenever this is not clearly impossible:

Theorem 1.2.2 (Dirichlet). *Let $q \geqslant 1$ be an integer and let a be an integer coprime with q. Then there exist infinitely many prime numbers p such that*

$$p \equiv a \,(\mathrm{mod}\, q).$$

In particular, taking $q = 10^k$ to be a power of 10, we can say that, whatever ending pattern of digits $\boldsymbol{d} = d_1 d_2 \cdots d_{k-1}$ we select, with $d_i \in \{0,1,2,3,4,5,6,7,8,9\}$, *provided* the last one d_{k-1} is not among $\{0,2,4,5,6,8\}$, there exist infinitely many prime numbers p with a decimal expansion where $\boldsymbol{d}$ are the final digits; for instance, taking $q = 1000$, $\boldsymbol{d} = 237$, we find

$$\begin{aligned}
&1237,\ 2237,\ 5237,\ 7237,\ 8237,\ 19237,\ 25237,\ 26237,\ 31237,\ 32237,\\
&\qquad 38237,\ 40237,\ 43237,\ 46237,\ 47237,\ 52237,\ 56237,\ 58237,\ 64237,\\
&\qquad\quad 70237,\ 71237,\ 73237,\ 77237,\ 82237,\ 85237,\ 88237,\ 89237,\ 91237,\ 92237
\end{aligned}$$

to be those prime numbers ending with 237 which are $\leqslant 100000$.

We will present the idea of the proof of this theorem in Chapter 4. As we will see, a crucial ingredient (but not the only one) is the simplest type of representation theory: that of groups which are both finite and commutative. In some sense, there is no better example to guess the power of representation theory than to see how even the simplest instance leads to such remarkable results.

Example 1.2.3 (The hydrogen atom). According to current knowledge, about 75% of the observable weight of the universe is accounted for by hydrogen atoms. In quantum mechanics, the possible states of an (isolated) hydrogen atom are described in terms of combinations of "pure" states, and the latter are determined by *discrete* data, traditionally called "quantum numbers" – so that the possible energy values of the system, for instance, form a discrete set of numbers, rather than a continuous interval.

Precisely, in the non-relativistic theory, there are four quantum numbers for a given pure state, denoted (n, ℓ, m, s) – "principal", "angular momentum", "magnetic" and "spin" are their usual names – which are all integers, except for s, with the restrictions

$$n \geqslant 1, \quad 0 \leqslant \ell \leqslant n-1, \quad -\ell \leqslant m \leqslant \ell, \quad s \in \{\pm 1/2\}.$$

It is rather striking that this simplest quantum-mechanical model of the hydrogen atom can be "explained" qualitatively by an analysis of the representation theory of the underlying symmetry group (see [**44**] or [**38**]), leading in particular to a natural explanation of the intricate structure of these four quantum numbers!

We will attempt to explain this in Section 6.4.

Example 1.2.4 ("Word" problems). For a prime number p, consider the finite group $\mathrm{SL}_2(\mathbf{F}_p)$ of 2×2 matrices of determinant 1 with coefficients in the finite field $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z}$. This group is generated by the two elements

$$s_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \qquad s_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \tag{1.1}$$

(this is a fairly easy fact from elementary group theory, see, e.g., [**33**, Th. 8.8] for $K = \mathbf{F}_p$, noting that s_1 itself generates the set of all upper-triangular transvections and s_2 the lower-triangular ones). Certainly the group is also generated by the elements of the set $S = \{s_1, s_1^{-1}, s_2, s_2^{-1}\}$, and in particular, for any $g \in \mathrm{SL}_2(\mathbf{F}_p)$, there exists an integer $k \geqslant 1$ and elements s_1, ..., s_k, each of which belongs to S, such that

$$g = s_1 \cdots s_k.$$

One may ask, how large can k be, at worse? The following result gives an answer:

Theorem 1.2.5 (Selberg, Brooks, Burger). *There exists a constant $C \geqslant 0$, independent of p and g, such that, with notation as above, we have*

$$k \leqslant C \log p$$

in all cases.

All proofs of this result depend crucially on ideas of representation theory, among other tools. And while it may seem to be rather simple and not particularly worth notice, the following *open* question should suggest that there is something very subtle here:

Question. Find an *algorithm* that, given p and $g \in \mathrm{SL}_2(\mathbf{F}_p)$, explicitly gives $k \leqslant C \log p$ and a sequence $(g_1, \ldots, g_k)$ such that

$$g = g_1 \cdots g_k.$$

For instance, how would you do with

$$g = \begin{pmatrix} 1 & (p-1)/2 \\ 0 & 1 \end{pmatrix}$$

(for $p \geqslant 3$)? Of course, one can take $k = (p-1)/2$ and $g_i = s_1$ for all i, but when p is large, this is much larger than what the theorem claims to be possible!

We will not prove Theorem 1.2.5, nor really say much more about the known proofs. However, in Section 4.7.1, we present more elementary results of Gowers [**16**] (and

Nikolov–Pyber [**30**]) which are much in the same spirit, and use the same crucial ingredient concerning representations of $\mathrm{SL}_2(\mathbf{F}_p)$.

In these three first examples, it turns out that representation theory appears in a similar manner: it is used to analyze functions on a group, in a way which is close to the theory of Fourier series or Fourier integrals – indeed, both of these can also be understood in terms of representation theory, as we will see, for the groups $(\mathbf{R}/\mathbf{Z})$ and $\mathbf{R}$, respectively. The next motivating example is purely algebraic:

Example 1.2.6 (Burnside's theorem on finite groups). Recall that if G is a finite group, it is called *solvable* if there are normal subgroups

$$1 \lhd G_k \lhd G_{k-1} \lhd \cdots \lhd G_1 \lhd G = G_0$$

where each successive quotient G_k/G_{k+1} is an abelian group.

Theorem 1.2.7 (Burnside). *Let G be a finite group. If the order of G is divisible by at most two distinct prime numbers, then G is solvable.*

This beautiful result is sharp in some sense: it is well-known that the group $\mathfrak{S}_5$ or its subgroup A_5 of index two are *not* solvable, and since its order 120 is divisible only by the primes 2, 3 and 5, we see that the analogue statement with 2 prime factors replaced with 3 is not true. (Also it is clear that the converse is not true either: any abelian group is solvable, and there are such groups of any order.)

This theorem of Burnside will be proved using representation theory of finite groups in Section 4.7.2 of Chapter 4, in much the same way as Burnside proceeded in the early 20th century. It is only in the late 1960's that a purely algebraic proof – not using representation theory – was constructed, first by Goldschmidt when the primes p and q are odd, and then by Bender and Matsuyama independently for the general case. There is a full account of this in [**20**, §7D], and although it is not altogether overwhelming in length, the reader who compares will probably agree that the proof based on representation theory is significantly easier to digest...

Remark 1.2.8. There are even more striking results, which are much more difficult; for instance, the famous "odd-order theorem" of Feit and Thompson states that if G has *odd* order, then G is necessarily solvable.

1.3. Prerequisites and notation

In Chapters 2 and 4, we depend only on the content of a basic graduate course in algebra: basic group theory, abstract linear algebra over fields, polynomial rings, finite fields, modules over rings, bilinear forms, and the tensor product and its variants. In other chapters, other structures are involved: groups are considered with a topology, measure spaces and integration theory occur, as well as basic Hilbert space theory and functional analysis. All these are considered at the level of introductory graduate courses.

Concerning algebra, we will use the following notation:

(1) For a set X, $|X| \in [0, +\infty]$ denotes its cardinal, with $|X| = \infty$ if X is infinite. There is no distinction in this text between the various infinite cardinals.
(2) Given a ring A, with a unit $1 \in A$, and A-modules M and N, we denote by $\mathrm{Hom}(M, N)$ or $\mathrm{Hom}_A(M, N)$ the space of A-linear maps from M to N.
(3) If E is a vector space over a field k, E' denotes the dual space $\mathrm{Hom}_k(E, k)$. We often use the duality bracket notation for evaluating linear maps on vectors, i.e.,

for $v \in E$ and $\lambda \in E'$, we write

$$\langle \lambda, v \rangle = \lambda(v).$$

(4) For $f : M \to N$ a map of A-modules, $\ker(f)$ and $\mathrm{Im}(f)$ denote the kernel and the image of f respectively.

(5) Given A and M, N as above, $M \otimes N$ or $M \otimes_A N$ denotes the tensor product of M and N. Recall that $M \otimes N$ can be characterized up to isomorphism by the existence of canonical isomorphisms

$$\mathrm{Hom}_A(M \otimes N, N_1) \simeq \mathrm{Bil}(M \times N, N_1)$$

for any A-module N_1, where the right-hand side is the A-module of all A-bilinear maps

$$\beta : M \times N \to N_1.$$

In particular, there is a bilinear map

$$\beta_0 : M \times N \longrightarrow M \otimes N$$

corresponding to $N_1 = M \otimes N$ and to the identity map in $\mathrm{Hom}_A(M \otimes N, N_1)$. One writes $v \otimes w$ instead of $\beta_0(v, w)$.

The elements of the type $v \otimes w$ in $M \otimes N$ are called *pure tensors*. Note that, usually, not all elements in the tensor product are pure tensors and that one can have $v \otimes w = v' \otimes w'$ even if $(v, w) \neq (v', w')$.

If $A = k$ is a field, and (e_i), (f_j) are bases of the k-vector spaces M and N, respectively, then $(e_i \otimes f_j)$ is a basis of $M \otimes N$. Moreover, any $v \in M \otimes N$ has then a *unique* expression

$$v = \sum_j v_j \otimes f_j$$

with $v_j \in M$ for all j.

(6) Given a ring A and A-modules given with linear maps

$$M' \xrightarrow{f} M \xrightarrow{g} M'$$

the composition is *exact* if $\mathrm{Im}(f) = \ker(g)$ in M.

(7) In particular, a sequence

$$0 \longrightarrow M' \xrightarrow{f} M$$

is exact if and only if $\ker(f) = 0$, which means that f is injective, and a sequence

$$M \xrightarrow{g} M'' \longrightarrow 0$$

is exact if and only if $\mathrm{Im}(g) = \ker(0) = M''$, i.e., if and only if g is surjective.

(8) A sequence

$$0 \longrightarrow M' \xrightarrow{f} M \xrightarrow{g} M'' \longrightarrow 0$$

where all three intermediate 3-terms compositions are exact is called a *short exact sequence.*

(9) Given a group G, we denote by $[G, G]$ the *commutator group* of G, which is generated by all commutators $[g, h] = ghg^{-1}h^{-1}$ (note that not all elements of $[G, G]$ are themselves commutators!). The subgroup $[G, G]$ is normal in G, and the quotient group $G/[G, G]$ is abelian; it is called the *abelianization* of G.

(10) We denote by $\mathbf{F}_p$ the finite field $\mathbf{Z}/p\mathbf{Z}$, for p prime, and more generally by $\mathbf{F}_q$ a finite field with q elements, where $q = p^n$, $n \geqslant 1$, is a power of p. In Chapter 4, we need some simple facts about these, in particular the fact that for each $n \geqslant 1$, there is – up to isomorphism – a unique extension $k/\mathbf{F}_p$ of degree n, i.e., a finite field k of order $q = p^n$. An element $x \in k$ is in $\mathbf{F}_p$ if and only if $x^p = x$ (e.g., because the equation $X^p - X = 0$ has at most p roots, and all $x \in \mathbf{F}_p$ are roots), and the group homomorphism

$$N \left\{ \begin{array}{ccc} k^\times & \longrightarrow & \mathbf{F}_p^\times \\ x & \mapsto & \prod_{j=0}^{n-1} x^{p^j} \end{array} \right.$$

(called the *norm* from k to $\mathbf{F}_p$) is surjective (e.g., because the kernel has at most $1 + p + p + \cdots + p^{n-1} = (p^n - 1)/(p - 1)$ elements, so the image has at least $p - 1$ elements). Moreover, the kernel of the norm is the set of all x which can be written as y/y^p for some $y \in k^\times$.

When considering a normed vector space E, we usually denote the norm by $\|v\|$, and sometimes write $\|v\|_E$, when more than one space (or norm) are considered simultaneously. When considering a Hilbert space H, we speak synonymously of inner product or positive-definite hermitian forms, denoted $\langle \cdot, \cdot \rangle$ or $\langle \cdot, \cdot \rangle_H$ if more than one space might be understood. We use the convention that a hermitian form is linear in the first variable, and conjugate-linear in the other, i.e., we have

$$\langle \alpha v, w \rangle = \alpha \langle v, w \rangle, \qquad \langle v, \alpha w \rangle = \bar{\alpha} \langle v, w \rangle,$$

for two vectors v, w and a scalar $\alpha \in \mathbf{C}$.

Concerning topology and measure theory, we recall the following definitions and facts. First, we always consider Hausdorff topological spaces, even if not explicitly mentioned. A *Borel measure* on a topological space X is a measure defined on the σ-algebra of Borel sets. A *Radon measure* is a Borel measure which is finite on compact subsets of X. The integral of a non-negative measurable function f, or of an integrable function f, with respect to μ, is denoted by either of the following

$$\int_X f(x) d\mu(x) = \int_X f d\mu.$$

CHAPTER 2

The language of representation theory

2.1. Basic language

We begin by restating formally the definition:

DEFINITION 2.1.1 (Linear representation). Let G be a group and let k be a field. A *linear representation of* G, defined over k, is a group homomorphism

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

where E is a k-vector space. The dimension of E is called the *degree*, or *rank* or simply *dimension* of ϱ, and is denoted $\deg(\varrho)$. It is also customary to just say that ϱ is a k-representation of G, and sometimes, when the homomorphism ϱ is clear from the context, one may say that E is a k-representation of G.

Given a representation ϱ acting on the vector space E, and an element $g \in G$, we usually write

$$\varrho(g)v$$

for the image of $v \in E$ under the linear transformation $\varrho(g)$. Such vectors are also called *G-translates* of v, or simply *translates*, when the context is clear. Similarly, when ϱ is clearly understood from context, one may simply write

$$gv = \varrho(g)v, \quad \text{or} \quad g \cdot v = \varrho(g)v.$$

The basic rules that ϱ satisfies are then the relations

$$\varrho(1_G)v = v \quad (gh)v = \varrho(gh)v = \varrho(g)(\varrho(h)v) = g(hv), \tag{2.1}$$

$$g^{-1}(gv) = \varrho(g^{-1})(\varrho(g)v) = v \tag{2.2}$$

for all g, $h \in G$ and $v \in E$, in addition to the linearity of $\varrho(g)$ for a given g.

This notation emphasizes the fact that ϱ is also the same as a left-action of the group G on the vector space E, the action being through linear maps (instead of arbitrary bijections of E). In this viewpoint, one thinks of ϱ as the equivalent data of the map

$$\begin{cases} G \times E & \longrightarrow & E \\ (g, v) & \mapsto & g \cdot v. \end{cases}$$

It should be already clear that representations exist in plenty – they are not among those mathematical objects that are characterized by their rarity. For instance, obviously, any subgroup G of $\mathrm{GL}(E)$ can be thought of as being given with a natural ("tautologous" is the adjective commonly used) representation

$$G \hookrightarrow \mathrm{GL}(E).$$

In a different style, for any group G and field k, we can form a vector space, denoted $k(G)$, with a basis $(e_g)_{g \in G}$ indexed by the elements of G (i.e., the k-vector space freely generated by the set G; if G is infinite, note that $k(G)$ is infinite-dimensional). Then we may let G act linearly on $k(G)$ by describing the transformation $\varrho_1(g)$ through its action on the basis vectors: we define

$$\varrho_1(g)e_h = e_{gh} \tag{2.3}$$

for all $g \in G$ and all basis vectors e_h. Then to check that ϱ_1 is a linear representation of G on E, it is enough to check (2.1). This is a simple exercise – we give details merely for completeness, but readers should attempt to perform this check, at least in a first reading. First, it is clear that $\varrho(1_G)$ acts as identity on the basis vectors, and hence is the identity transformation. Now, given g_1, $g_2 \in G$ and a basis vector e_h, its image under $\varrho_1(g_1g_2)$ is $e_{g_1g_2h}$ by definition. And since $\varrho_1(g_2)e_h$ is the basis vector e_{g_2h}, we also have

$$\varrho_1(g_1)(\varrho_1(g_2)e_h) = e_{g_1g_2h} = \varrho_1(g_1g_2)e_h$$

which, h being arbitrary, means that $\varrho_1(g_1g_2) = \varrho_1(g_1)\varrho_1(g_2)$. By taking $g_2 = g_1^{-1}$ this confirms that ϱ_1 is a homomorphism into $\mathrm{GL}(k(G))$.

Another easily defined representation is the *right-regular representation*, or simply *regular* representation (over k): let[1] $C_k(G)$ be the space of all functions

$$f\,:\,G \to k$$

(with pointwise addition and scalar multiplication of functions; we will often write $C(G)$ for $C_k(G)$ when the field is clear in context). One defines $\mathrm{reg}(g)$ acting on $C_k(G)$ by the rule

$$\mathrm{reg}(g)f(x) = f(xg)$$

for all $f \in C_k(G)$, $g \in G$, where $x \in G$ is the point at which the new function $\mathrm{reg}(g)f \in C_k(G)$ is evaluated.[2] It is again a simple matter – that the reader should attempt, if only because the order of evaluation might seem to be wrong! – to check that reg is a representation: for $f \in E$, g, $h \in G$, we get that $\mathrm{reg}(gh)f$ maps x to

$$\mathrm{reg}(gh)f(x) = f(xgh),$$

while, $\mathrm{reg}(h)f$ being the function $f_1\,:\,y \mapsto f(yh)$, we see that $\mathrm{reg}(g)\,\mathrm{reg}(h)f = \mathrm{reg}(g)f_1$ maps x to

$$f_1(xg) = f((xg)h) = f(xgh),$$

which completes the check that $\mathrm{reg}(gh) = \mathrm{reg}(g)\,\mathrm{reg}(h)$.

In all three examples, the representation map ϱ is injective (it is clear in the second one and easily checked in the third). This is certainly not always the case: indeed, for any group G and field k, there is also a “trivial” representation of G of degree 1 defined over k, which simply maps every $g \in G$ to $1 \in k^\times = \mathrm{GL}(k)$. Obviously, this is not injective unless $G = 1$. Note that one shouldn't dismiss this trivial representation as obviously uninteresting: as we will see quite soon, it does have an important role to play.

Still we record the names of these two types of representations:

Definition 2.1.2 (Faithful and trivial representations). Let G be a group and let k be a field.

(1) A representation ϱ of G defined over k is *faithful* if ϱ is injective, i.e., if $\ker(\varrho) = \{1\}$ in G.

(2) A representation ϱ of G on a k-vector space E is *trivial* if $\varrho(g) = 1$ is the identity map of E for all $g \in G$, i.e., if $\ker(\varrho) = G$.

[1] The notation is not completely standard.

[2] There is also a *left-regular representation*, where $\mathrm{reg}_{left}(g)f(x) = f(g^{-1}x)$.

Remark 2.1.3. Sometimes only the representation of degree 1 (with $E = k$) mapping g to $1 \in k^\times$ is called “the” trivial representation. We will denote by $\mathbf{1}$ this one-dimensional representation (when G and k are clear in context, or $\mathbf{1}_G$ if only k is).

These examples are extremely general. Before continuing, here are others which are extremely specific – but still very important. We take $k = \mathbf{C}$; then we have the exponential $z \mapsto e^z$, which is a group homomorphism from $(\mathbf{C}, +)$ to $(\mathbf{C}^\times, \cdot)$, or in other words, to $\mathrm{GL}_1(\mathbf{C}) = \mathrm{GL}(\mathbf{C})$. This means the exponential is a 1-dimensional representation (over $\mathbf{C}$) of the additive group of the complex numbers. One can find variants:

- If $G = \mathbf{R}$ or $\mathbf{C}$, then for any $s \in \mathbf{C}$, the map

$$\chi_s \,:\, x \mapsto e^{sx} \tag{2.4}$$

is a one-dimensional representation.

- If $G = \mathbf{R}/\mathbf{Z}$, then for any $m \in \mathbf{Z}$, the map

$$e_m \,:\, x \mapsto e^{2i\pi mx} \tag{2.5}$$

is a one-dimensional representation of G (one must check that this is well-defined on $\mathbf{R}/\mathbf{Z}$, but this is the case since $e^{2i\pi mn} = 1$ for any $n \in \mathbf{Z}$; indeed, no other representation χ_s of $\mathbf{R}$, for $s \in \mathbf{C}$, has this property since $\chi_s(1) = 1$ means $e^s = 1$.)

- If $q \geqslant 1$ is an integer and $G = \mathbf{Z}/q\mathbf{Z}$ if the additive group of integers modulo q, then for any $m \in \mathbf{Z}/q\mathbf{Z}$, the map

$$x \mapsto e^{2i\pi mx/q} \tag{2.6}$$

is well-defined and it is a one-dimensional representation of G. Indeed, note that $e^{2i\pi\tilde{m}x/q}$ is independent of the choice of a representative $\tilde{m} \in \mathbf{Z}$ of $m \in \mathbf{Z}/q\mathbf{Z}$, since replacing $\tilde{m}$ by $\tilde{m} + kq$ just multiplies the value by $e^{2i\pi xk} = 1$.

More examples, many of which are defined without the intermediate results and language, can be found in Section 2.6, and some readers may want to read that section first (or at least partly) to have some more concrete examples in mind.

Although one can thus see that there are “many” representations in a certain sense, as soon as we try to “compare” them, the impression emerges that this abundance is – for given G and field k – of the same type as the abundance of vector spaces (in contrast with, for instance, the similarly striking abundance of k-algebras): although they may arise in every corner, many of them are actually the same. In other words, quite often, the representations of G over k can be classified in a useful way. To go into this, we must explain how to relate possibly different representations.

Definition 2.1.4 (Morphism of representations). Let G be a group and let k be a field. A *morphism*, or homomorphism, between representations ϱ_1 and ϱ_2 of G, both defined over k and acting on the vector spaces E_1 and E_2, respectively, is a k-linear map

$$\Phi \,:\, E_1 \longrightarrow E_2$$

such that

$$\Phi(\varrho_1(g)v) = \varrho_2(g)(\Phi(v)) \in E_2,$$

for all $g \in G$ and $v \in E_1$. One also says that Φ *intertwines* ϱ_1 and ϱ_2, or is an *an intertwining operator* between them, and one may denote this by $\varrho_1 \xrightarrow{\Phi} \varrho_2$.

This definition is also better visualized as saying that, for all $g \in G$, the square diagram

$$\begin{tikzcd} E_1 \arrow[r, "\Phi"] \arrow[d, "\varrho_1(g)"'] & E_2 \arrow[d, "\varrho_2(g)"] \\ E_1 \arrow[r, "\Phi"] & E_2 \end{tikzcd}$$

of linear maps commutes, or – even more concisely – by omitting the mention of the representations and writing

$$\Phi(g \cdot v) = g \cdot \Phi(v)$$

for $g \in G$, $v \in E_1$.

It is also easy to see that the set of homomorphisms from ϱ_1 to ϱ_2, as representations of G, is a k-vector subspace of $\mathrm{Hom}(E_1, E_2)$, which we denote $\mathrm{Hom}_G(\varrho_1, \varrho_2)$. (As we will see, this vector space may be, and is often, reduced to 0!).

The following is clear, but also of crucial importance:

Proposition 2.1.5 (Functoriality). *Let G be a group and k a field.*

(1) *For any representation ϱ of G and a vector space E, the identity map on E is a homomorphism $\varrho \longrightarrow \varrho$.*

(2) *Given representations ϱ_1, ϱ_2 and ϱ_3 on E_1, E_2 and E_3 respectively, and morphisms*

$$E_1 \xrightarrow{\Phi_1} E_2 \xrightarrow{\Phi_2} E_3,$$

the composite $E_1 \xrightarrow{\Phi_2 \circ \Phi_1} E_3$ is a morphism between ϱ_1 and ϱ_3.

Remark 2.1.6 (The category of representations). In the language of category theory (which we will only use incidentally in remarks in these notes), this proposition states that the representations of a given group G over a given field k are the objects of a *category* with morphisms given by the intertwining linear maps.

If a morphism Φ is a bijective linear map, its inverse Φ^{-1} is also a morphism (between ϱ_2 and ϱ_1), and it is therefore justified to call Φ an *isomorphism* between ϱ_1 and ϱ_2. Indeed, using the diagram above, we find that the relation

$$\varrho_2(g) \circ \Phi = \Phi \circ \varrho_1(g)$$

is equivalent in that case to

$$\Phi^{-1} \circ \varrho_2(g) = \varrho_1(g) \circ \Phi^{-1},$$

which is the desired fact that Φ^{-1} be an intertwining operator between ϱ_2 and ϱ_1.

As an example, if a vector subspace $F \subset E$ is stable under all operators $\varrho(g)$ (i.e., $\varrho(g)(F) \subset F$ for all $g \in G$), then the restriction of $\varrho(g)$ to F defines a homomorphism

$$\tilde{\varrho} : G \longrightarrow \mathrm{GL}(F),$$

which is therefore a k-representation of G, and the inclusion linear map

$$i : F \hookrightarrow E$$

is a morphism of representations. One speaks, naturally, of a *subrepresentation* of ϱ or, if the action is clear in context, of E itself.

Example 2.1.7 (Trivial subrepresentations). Consider the special case where F is the space of all $v \in E$ which are pointwise invariant under G: $v \in F$ if and only if

$$g \cdot v = v \text{ for all } g \in G.$$

Because G acts by linear maps on E, this subspace F, also denoted $F = E^G$, is a linear subspace of E and a subrepresentation of ϱ. Note that the representation of G on E^G is trivial, in the sense of Definition 2.1.2. This means that if n is the dimension[3] of E^G, and if $\mathbf{1}^n = k^n$ denotes the k-vector space of dimension n with a trivial action of G, we have an isomorphism

$$\mathbf{1}^n \xrightarrow{\sim} E^G$$

(by fixing any basis of E^G). Of course, it is possible – and is frequently the case – that $E^G = 0$.

This space of invariants is the largest subrepresentation of E (for inclusion) which is trivial. More individually, any non-zero vector $v \in E$ which is invariant under G defines a trivial subrepresentation of dimension 1, i.e., an injective morphism

$$\begin{cases} \mathbf{1} & \hookrightarrow & E \\ t & \mapsto & tv \end{cases}$$

of representations. This gives a k-linear isomorphism

$$E^G \simeq \mathrm{Hom}_G(\mathbf{1}, E) \tag{2.7}$$

(the reciprocal map sending $\Phi : \mathbf{1} \to E$ to $\Phi(1)$).

Because fixed points or invariant vectors of various kinds are often of great importance, we see here how useful the trivial representation can be. To give a simple – but, it turns out, very useful example! – the invariant subspace of the regular representation is the one-dimensional subspace of constant (k-valued) functions on G: if $\varphi \in C_k(G)^G$, we have

$$\varphi(x) = \mathrm{reg}(g)\varphi(x) = \varphi(xg)$$

for all x and g, and taking $x = 1$ shows that φ is constant.

On the other hand, note that $k(G)^G$ is *zero* if G is infinite, and one-dimensional, generated by

$$\sum_{g \in G} e_g \in k(G)$$

if G is finite.

Example 2.1.8 (Invariants under normal subgroups). Consider again a k-representation ϱ of G, acting on E. The space E^G of invariants is a subrepresentation, obviously trivial, as in the previous example. A very useful fact is that if we take the vectors invariant under a subgroup of G, provided it is a normal subgroup, we still obtain a subrepresentation of E, though not a trivial one usually.

Lemma 2.1.9 (Invariants under normal subgroups). *Let k be a field, let G be a group and $H \triangleleft G$ a normal subgroup. Then for any k-representation ϱ of G acting on E, the subspace*

$$E^H = \{v \in E \mid \varrho(h)v = v \text{ for all } h \in H\}$$

is a subrepresentation of ϱ.

Proof. Let $v \in E^H$ and $g \in G$. We want to check that $w = \varrho(g)v \in E^H$, and for this we pick $h \in H$ and we write simply

$$\varrho(h)w = \varrho(hg)v = \varrho(g)\varrho(g^{-1}hg)v,$$

and since $h' = g^{-1}hg$ is in H (because H is normal by assumption) and $v \in E^H$, we get $\varrho(h)w = \varrho(g)v = w$ as desired. □

[3] Which may be finite or infinite.

The reader should look for examples where H is not normal and E^H is not stable under the action of G, as well as for examples where E^H is not a trivial representation of G.

Example 2.1.10 (Regular representation). Consider the two examples of representations ϱ_1 and reg associated to a group G and field k that were discussed just after the Definition 2.1.1. We claim that ϱ_1 (acting on $k(G)$) is isomorphic to a subrepresentation of reg (acting on $C(G)$). To see this, we define $\Phi\,:\,k(G)\rightarrow C(G)$ by mapping a basis vector e_g, $g\in G$, to the characteristic function of the single point g^{-1}:

$$\Phi(e_g)\,:\,x\mapsto\begin{cases}1 & \text{if } x=g^{-1},\\ 0 & \text{otherwise.}\end{cases}$$

The linear map defined in this way is injective – indeed, $\Phi(v)$ is the function mapping $g\in G$ to the coefficient of the basis element $e_{g^{-1}}$ in the expression of v, and can only be identically zero if v is itself 0 in $k(G)$. We check now that Φ is a morphism of representations. In $k(G)$, we have $g\cdot e_h=e_{gh}$, and in $C(G)$, we find that $g\cdot\Phi(e_h)=\mathrm{reg}(g)\Phi(e_h)$ maps x to

$$\Phi(e_h)(xg)=\begin{cases}1 & \text{if } xg=h^{-1}\text{ or } x=h^{-1}g^{-1}=(gh)^{-1},\\ 0 & \text{otherwise.}\end{cases}$$

which precisely says that

$$\Phi(g\cdot e_h)=g\cdot\Phi(e_h).$$

The map Φ is an isomorphism if G is finite, but not otherwise; indeed, the image $\mathrm{Im}(\Phi)$ is always equal to the subspace of functions which are zero except at finitely many points.

Remark 2.1.11. The last example makes it fairly clear that our basic definitions will require some adaptations when infinite groups are considered. Typically, if G has a topological structure – compatible with the group operation – the regular representation will be restricted to functions with a certain amount of smoothness or regularity. We will come back to this in Chapter 3 (and later).

We will now discuss the basic formalism of representation theory – roughly speaking, how to manipulate some given representation or representations to obtain new ones. This involves different aspects, as one may try to operate at the level of the vector space E, or of the group G, or even of the field k. The latter is of less importance in these notes, where k will be $\mathbf{C}$ most of the time after this chapter, but we will mention it briefly nevertheless. The other two are, however, of fundamental importance.

2.2. Formalism: changing the space

This part of the formalism is the most straightforward. The basic philosophy is simply that essentially any operation of linear or multilinear algebra can be performed on a space E on which a group G acts in such a way that G has a natural action on the resulting space. This particularly transparent when interpreting representations of G as modules over the group algebra, as explained in Chapter 3, but we will present the basic examples from scratch. However, before reading further, we suggest to the reader that she try to come up with the definition of the following objects (where the field k and the group G are always fixed):

– Quotients of representations, sum and intersection of subrepresentations;

– The kernel and image of a morphism of representations;
– Exact sequences, in particular, short exact sequences, of representations;
– The direct sum of representations;
– The tensor product of two representations;
– The symmetric powers or alternating powers of a representation;
– Given a representation ϱ acting on E, the dual (also called contragredient) of ϱ acting on the linear dual space $E' = \mathrm{Hom}_k(E, K)$, and the associated representation of G acting on the space of k-linear maps $\mathrm{End}_k(E) = \mathrm{Hom}_k(E, E)$.

As will be seen, only the last one may be not entirely obvious, and this is because there are in fact two possible answers (though, as we will explain, one of them is much more interesting and important).

Here is an abstract presentation of the mechanism at work; although we will give full details in each case, it is also useful to see that a single process is at work.

Proposition 2.2.1 (Functorial representations). *Let k be a field and G a group. Let T be any functor on the category of k-vector spaces, i.e., any rule assigning a vector space $T(E)$ to any k-vector space E, and a map*

$$T(f) : T(E_1) \to T(E_2)$$

to any linear map $f : E_1 \to E_2$, with the properties that

$$\begin{cases} T(f \circ g) = T(f) \circ T(g), \\ T(1_E) = 1_{T(E)}. \end{cases} \tag{2.8}$$

Then given a k-representation

$$\varrho : G \longrightarrow \mathrm{GL}(E),$$

the vector space $T(E)$ has a linear action

$$\pi = T(\varrho) : G \longrightarrow \mathrm{GL}(T(E))$$

given by

$$\pi(g) = T(\varrho(g)).$$

Moreover, for any homomorphism $\varrho_1 \xrightarrow{\Phi} \varrho_2$ of representations of G, the k-linear map $T(\Phi)$ is a homomorphism $T(\varrho_1) \longrightarrow T(\varrho_2)$, and this construction is compatible with composition and identity. In particular, $T(\varrho)$ depends, up to isomorphism of representations, only on the isomorphism class of ϱ itself.

This is a direct translation of the "functoriality" property of morphisms of representations noted in Proposition 2.1.5.

2.2.1. Quotients, kernels, images,... We have defined subrepresentations already. The operation of sum and intersection of subspaces, when applied to subrepresentations, lead to other subrepresentations – this should be clear.

Quotients are equally natural objects to consider. Given a representation ϱ of G on E, and a subspace $F \subset E$ which is a *subrepresentation* of E, or in other words, such that $\varrho(g)$ always leaves F invariant, the quotient vector space $H = E/F$ also has a natural linear action of G, simply induced by ϱ: given $v \in H$ and $g \in G$, the action $g \cdot v$ is the image in H of $\varrho(g)\tilde{v}$ for any $\tilde{v} \in E$ mapping to v under the canonical projection map $E \to H$. This is well-defined because if $\tilde{v}_1$ is another such vector, we have $\tilde{v}_1 = \tilde{v} + w$ with $w \in F$, hence

$$\varrho(g)\tilde{v}_1 - \varrho(g)\tilde{v} = \varrho(g)w$$

also lies in F, and has image 0 in H.

Another global description of this action is that it is such that the projection map

$$E \longrightarrow H = E/F$$

is then a morphism of representations, just like the inclusion map $F \longrightarrow E$ is one.

In the same vein, given now a morphism

$$\Phi \,:\, E_1 \longrightarrow E_2$$

of k-representations of G, we can see that the standard vector spaces associated to Φ are all themselves representations of G:
– The kernel $\ker(\Phi) \subset E_1$ is a subrepresentation of E_1;
– The image $\mathrm{Im}(\Phi) \subset E_2$ is a subrepresentation of E_2;
– The natural linear isomorphism

$$E_1/\ker(\Phi) \simeq \mathrm{Im}(\Phi)$$

(induced by Φ) is an isomorphism of representations;
– The cokernel $\mathrm{coker}(\Phi) = E_2/\mathrm{Im}(\Phi)$ is a representation of G, as quotient of two representations.

These facts are consequences of the definitions, and specifically of the linearity of the actions of G.

2.2.2. Coinvariants. If we go back to Example 2.1.7, and in particular the identification (2.7) of the homomorphisms from $\mathbf{1}$ to a representation ϱ, one may ask if there is a similar description of the space

$$\mathrm{Hom}_G(E, \mathbf{1})$$

of homomorphisms *from* ϱ *to* the trivial one.

By definition, an element in this space is a k-linear form $E \xrightarrow{\lambda} k$ such that for all $v \in E$ and $g \in G$, we have

$$\lambda(g \cdot v) = \lambda(v).$$

In other words, λ is exactly the same as a linear form which factors through the subspace E_1 of E spanned by all vectors of the form

$$g \cdot v - v, \qquad g \in G,\ v \in E,$$

or equivalently it corresponds to a linear form

$$E/E_1 \longrightarrow k.$$

Note that E_1 is also a subrepresentation of ϱ, since

$$h \cdot (g \cdot v - v) = hg \cdot v - h \cdot v = (hgh^{-1})v_1 - v_1 \tag{2.9}$$

with $v_1 = h \cdot v$. Hence E/E_1 has an induced structure of representation of G. In fact, this action on E/E_1 is trivial, since $g \cdot v = v$ modulo E_1 for all g and v.

The space E/E_1 is called the space of *coinvariants* of ϱ, and is denoted E_G or ϱ_G. It is the "largest" quotient of ϱ which is a trivial representation of G (like the invariant space, it may well be zero) and by the above, we can write

$$\mathrm{Hom}_G(\varrho, \mathbf{1}) \simeq \mathrm{Hom}_k(\varrho_G, k),$$

which identifies the space of homomorphisms to the trivial representation with the linear dual vector space of the coinvariant space.

We leave to the reader the simple exercise of checking that the analogue of Lemma 2.1.9 holds for the coinvariants under a normal subgroup: if $H \lhd G$, the quotient E_H has an induced structure of representation of G (use (2.9)).

2.2.3. Direct sums, exact sequences, irreducibility and semisimplicity. The simplest operation that can be performed on representations is the direct sum. Given G and k, as usual, and k-representations ϱ_1, ϱ_2 of G on E_1 and E_2, respectively, the direct sum $\varrho_1 \oplus \varrho_2$ is the representation

$$G \longrightarrow \mathrm{GL}(E_1 \oplus E_2)$$

such that

$$g \cdot (v_1 + v_2) = \varrho_1(g)v_1 + \varrho_2(g)v_2,$$

for all $v = v_1 + v_2 \in E_1 \oplus E_2$, or more suggestively

$$g \cdot (v_1 + v_2) = g \cdot v_1 + g \cdot v_2.$$

By definition, we see that the subspaces E_1, E_2 or $E = E_1 \oplus E_2$ are subrepresentations of $\varrho_1 \oplus \varrho_2$, and that

$$(\varrho_1 \oplus \varrho_2)/\varrho_1 \simeq \varrho_2 \tag{2.10}$$

for instance (the corresponding isomorphism being induced by $v_1 + v_2 \mapsto v_2$).

One can consider more than two factors: for an arbitrary family $(\varrho_i)_{i \in I}$ of k-representations, with ϱ_i acting on E_i, one can define a representation of G on the direct sum

$$E = \bigoplus_{i \in I} E_i$$

by linearity again from the actions of G on each subspace E_i of E.

Note the general relations

$$\deg(\varrho_1 \oplus \varrho_2) = \deg(\varrho_1) + \deg(\varrho_2), \quad \deg(\bigoplus_{i \in I} \varrho_i) = \sum_{i \in I} \deg(\varrho_i)$$

with obvious conventions when the sum is infinite. Equally useful are the natural isomorphisms

$$\mathrm{Hom}_G(\varrho, \varrho_1 \oplus \varrho_2) \simeq \mathrm{Hom}_G(\varrho, \varrho_1) \oplus \mathrm{Hom}_G(\varrho, \varrho_2),$$
$$\mathrm{Hom}_G(\varrho_1 \oplus \varrho_2, \varrho) \simeq \mathrm{Hom}_G(\varrho_1, \varrho) \oplus \mathrm{Hom}_G(\varrho_2, \varrho),$$

and similarly for an arbitrary (finite) number of summands.[4]

Another generalization of the direct sum, based on (2.10), considers any representation ϱ of G acting on E, with an injection

$$\Phi \,:\, \varrho_1 \hookrightarrow \varrho$$

such that

$$\varrho/\varrho_1 = \varrho/\operatorname{Im}(\Phi) \simeq \varrho_2 \tag{2.11}$$

as k-representations. However, although there exists of course always a subspace $E_2 \subset E$ such that

$$E = \operatorname{Im}(\Phi) \oplus E_2 \simeq E_1 \oplus E_2$$

as k-vector spaces, *it is not always the case that E_2 can be found as a subrepresentation of ϱ.* When that happens, this subrepresentation on E_2 (say $\tilde{\varrho}_2$) is necessarily isomorphic

[4] Recall that $\mathrm{Hom}_k(\bigoplus E_i, E)$ is not isomorphic to the direct sums of the $\mathrm{Hom}_k(E_i, E)$ if the index set is infinite – e.g. for $E = k$, the dual of a direct sum is the product of the duals, which is different for infinitely many factors.

to ϱ_2 (since $\tilde{\varrho}_2 \simeq (\varrho_1 \oplus \tilde{\varrho}_2)/\varrho_1 \simeq \varrho/\varrho_1 \simeq \varrho_2$, as representations of G). A useful equivalent criterion for the existence of such a completementary subrepresentation is the following:

Lemma 2.2.2. *Let G be a group, k a vector space and $\varrho : G \longrightarrow \mathrm{GL}(E)$ a representation.*

(1) *If $E = E_1 \oplus E_2$ is a decomposition of E such that E_1 is a subrepresentation of ϱ, then E_2 is also one if and only if the linear projection map*

$$\Phi \begin{cases} E & \longrightarrow & E \\ v = v_1 + v_2 & \mapsto & v_1 \end{cases}, \qquad v_1 \in E_1, \ v_2 \in E_2,$$

with image E_1 and kernel E_2 is an intertwiner*, i.e., if $\Phi \in \mathrm{Hom}_G(E, E)$.*

(2) *If $E_1 \subset E$ is a subrepresentation, there exists a linear complement E_2 which is a subrepresentation if and only if there exists an intertwiner in $\mathrm{Hom}_G(E, E)$ which is a projection, and such that $\mathrm{Im}(\Phi) = E_1$. A stable complement E_2 is then given by $E_2 = \ker \Phi$.*

Proof. This is elementary, and (2) is of course a consequence of (1), which follows by noting first that if Φ is an intertwiner, the kernel $\ker \Phi = E_2$ is a subrepresentation, while conversely, if E_2 is a subrepresentation, we get from $v = v_1 + v_2$ with $v_i \in E_i$ the decompositions $\varrho(g)v = \varrho(g)v_1 + \varrho(g)v_2$ with $\varrho(g)v_i \in E_i$ again, and hence $\Phi(\varrho(g)v) = \varrho(g)v_2 = \varrho(g)\Phi(v)$. □

In certain circumstances, the existence of the subrepresentation complementary to ϱ_1 is always valid (for instance, for finite groups when k has characteristic 0, as we will discuss in Chapter 4). Here is an example where it fails: consider the additive group $G = \mathbf{Z}/p\mathbf{Z}$ and the field $k = \mathbf{Z}/p\mathbf{Z}$, and the representation

$$\varrho \begin{cases} G & \longrightarrow & \mathrm{GL}_2(k) \\ x & \mapsto & \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} \end{cases} \tag{2.12}$$

(we leave as an exercise to check, if needed, that this is a homomorphism; note however how its existence depends on the fact that G is a subgroup of the additive group of k). In terms of the canonical basis (e_1, e_2) of k^2, this means that

$$x \cdot (\alpha e_1 + \beta e_2) = (\alpha + x\beta)e_1 + \beta e_2.$$

Therefore the subspace $E_1 = ke_1$ is a subrepresentation of G, indeed, it is isomorphic to the trivial representation $\mathbf{1}_G$ since e_1 is invariant under the action of G (which is obvious when looking at the matrix representation). We claim that there is *no* subspace E_2 complementary to E_1. This can be checked by direct computations (taking a hypothetical basis vector $f = \alpha e_1 + \beta e_2$ of E_2), but also more abstractly by noting that the quotient representation ϱ/E_1 (note the slight abuse of notation, which is quite usual) is itself the trivial representation (this should be checked from the definition; in terms of the matrix representation, it amounts to the fact that the bottom-right coefficients of $\varrho(x)$ are all equal to 1). Thus if E_2 were to exist, we would have, by the above, an isomorphism

$$\varrho \simeq \mathbf{1}_G \oplus \mathbf{1}_G,$$

which is a trivial representation of dimension 2. Since ϱ is certainly not trivial, this would be a contradiction.

Coming back to the general case where (2.11) holds, it is often summarized, as in linear algebra, by a *short exact sequence*

$$0 \rightarrow \varrho_1 \longrightarrow \varrho \longrightarrow \varrho_2 \rightarrow 0.$$

When ϱ is isomorphic to the direct sum of ϱ_1 and ϱ_2, one says that the exact sequence *splits*. And more generally, a sequence of homomorphisms of k-representations of G is exact, if and only if, it is exact as a sequence of maps of k-vector spaces.

Any time a natural representation can be written (up to isomorphism) as a direct sum, or even an extension, of smaller representations, this gives very useful information on the representation. Typically one wishes to perform such decompositions as long as it is possible. The obvious limitation is that a representation ϱ might not have any non-trivial subrepresentation to try to "peel off". This leads to the following very important definitions:

Definition 2.2.3 (Irreducible, semisimple representations). Let G be a group and k a field.

(1) A k-representation ϱ of G acting on E is *irreducible* if and only if $E \neq 0$ and there is no subspace of E stable under ϱ, except 0 and E itself (in other words, if there is no subrepresentation of ϱ except 0 and ϱ itself).

(2) A k-representation ϱ of G is *semisimple* if it can be written as a direct sum of subrepresentations, each of which is irreducible:

$$\varrho \simeq \bigoplus_{i \in I} \varrho_i$$

for some index set I and some irreducible representations ϱ_i.

In a semisimple representation, note that some of the summands ϱ_i might be themselves isomorphic. On the other hand, we will see later that, up to permutation, the irreducible summands are uniquely determined by ϱ (up to isomorphism of representations, of course): this is part of the Jordan-Hölder-Noether Theorem 2.7.1.

Not all representations of a group are semisimple, but irreducible representations are still fundamental "building blocks" for representations in general. An essential feature of irreducible representations, which is formalized in Schur's Lemma 2.2.4, is that these "building blocks" are "incommensurable", in some sense: two non-isomorphic irreducible representations can have "no interaction".

Lemma 2.2.4 (Schur's Lemma, I). *Let G be a group and let k be a field.*

(1) *Given an* irreducible *k-representation π of G and an arbitrary representation ϱ of G, any G-homomorphism $\pi \longrightarrow \varrho$ is either 0 or injective, and any G-homomorphism $\varrho \longrightarrow \pi$ is either 0 or surjective.*

(2) *Given irreducible k-representations π and ϱ of G, a homomorphism $\pi \longrightarrow \varrho$ is either 0 or is an isomorphism; in particular, if π and ϱ are not isomorphic, we have*

$$\mathrm{Hom}_G(\pi, \varrho) = 0.$$

Proof. (1) Given a morphism Φ from π to ϱ, we know that its kernel is a subrepresentation of π; but if π is irreducible, the only possibilities are that the kernel be 0 (then Φ is injective) or that it is π itself (then Φ is 0). Similarly for a morphism from ϱ to π, the image is either 0 or π itself.

(2) From (1), if Φ is non-zero and has irreducible source and target, it must be an isomorphism. (Recalling that, by definition, an irreducible representation is non-zero, we see that these are exclusive alternatives.) □

Although an arbitrary representation of a group may fail to contain irreducible subrepresentations, we can always find one in a finite-dimensional non-zero representation, by simply selecting a non-zero subrepresentation of minimal dimension. Hence:

Lemma 2.2.5 (Existence of irreducible subrepresentations). *Let G be a group, k a field and ϱ a non-zero k-representation of G. If ϱ is finite-dimensional, there exists at least one irreducible subrepresentation of G contained in ϱ.*

Remark 2.2.6 (Cyclic vector). It is tempting to suggest a more general argument by saying that, given a non-zero representation $G \longrightarrow \mathrm{GL}(E)$, and given $v \neq 0$, the linear span of the vectors $\varrho(g)v$, $g \in G$ should be irreducible – it is after all the smallest subrepresentation of G containing v for inclusion (indeed, any $F \subset E$ which is stable under the action of G and contains v must contain all such vectors, hence also their linear span). However, in general, this space is *not* irreducible.

For instance, consider the group $G = \mathbf{Z}/p\mathbf{Z}$ with $p \geqslant 3$ prime, and the representation on $\mathbf{C}^2$ given by

$$x \cdot (z_1, z_2) = (e^{2i\pi x/p} z_1, e^{-2i\pi x/p} z_2).$$

Since the two axes are invariant under this action, it is of course not irreducible. However, taking $v = (1, 1) \in \mathbf{C}^2$, we see that the span of all $x \cdot v$ contains $(1, 1)$ and

$$1 \cdot (1, 1) = (e^{2i\pi/p}, e^{-2i\pi/p}),$$

and since

$$\begin{vmatrix} 1 & 1 \\ e^{2i\pi/p} & e^{-2i\pi/p} \end{vmatrix} = -2i \sin\Bigl(\frac{2\pi}{p}\Bigr) \neq 0,$$

this vector does "generate" the whole space.

For a given representation $\varrho : G \longrightarrow \mathrm{GL}(E)$ of a group G, if there exists a non-zero vector $v \in E$ such that its translates span E, it is customary to say that ϱ is a *cyclic representation* and that v is then a *cyclic vector* (which is far from unique usually). For a given vector v, the space generated by the vectors $\varrho(g)v$, which is a cyclic subrepresentation of ϱ, is called the representation *generated* by v.

The example above generalizes to any group G and any representation of the type

$$\varrho = \bigoplus_{1 \leqslant i \leqslant k} \varrho_i$$

where the ϱ_i are pairwise non-isomorphic irreducible representations of G: taking $v = (v_i)$ in the space of ϱ, where each v_i is non-zero, it follows from the linear independence of matrix coefficients (Theorem 2.7.24 below) that v is a cyclic vector for ϱ.

The simplest examples of irreducible k-representations of G are the 1-dimensional representations

$$\chi : G \longrightarrow \mathrm{GL}(k) \simeq k^{\times}$$

(sometimes called the *characters* of G, though this clashes with another more general notion of character, as seen below in Definition 2.7.31) since there is no possible intermediate subrepresentations here! If we see characters as $k^{\times}$-valued functions on G, then we see quickly that they define isomorphic representations if and only if the functions are equal.

In particular, the trivial representation $\mathbf{1}_G$ is irreducible (and it may be the only 1-dimensional representation of G). Thus also any trivial representation on a vector space E is semisimple, since it can be written as a direct sum of trivial one-dimensional subrepresentations

$$E \simeq \bigoplus_{i \in I} k e_i,$$

after choosing a basis $(e_i)_{i\in I}$ of E. This shows, in passing, that the decomposition of a semisimple representation as a sum of irreducible ones is usually not unique, just as the choice of a basis of a vector space is not unique.

On the other hand, the 2-dimensional representation in (2.12) is not semisimple (this is intuitively clear, even if it might need some ad-hoc argument to check at this point; from the Jordan-Hölder-Noether Theorem below, this follows because if it were semisimple, it would have to be trivial, which it is not.)

The following lemma is also very useful as it shows that semisimple representations are stable under the operations we have already seen:

Lemma 2.2.7 (Stability of semisimplicity). *Let G be a group and let k be a field. If ϱ is a semisimple k-representation of G, then any subrepresentation of ϱ is also semisimple, and any quotient representation of ϱ is also semisimple.*

One should be careful that, if ϱ acts on E and we have stable subspaces E_i such that G acts on E_i via ϱ_i and

$$E = \bigoplus_{i\in I} E_i,$$

it does *not* follow that any subrepresentation is of the type

$$\bigoplus_{i\in J} E_i$$

for some $J \subset I$. This is false even for G trivial, where the only irreducible representation is the trivial one, and writing a decomposition of E amounts to choosing a basis. Then there are usually many subspaces of E which are *not* literally direct sums of a subset of the basis directions (e.g., $E = k \oplus k$ and $F = \{(x, x) \mid x \in k\}$).

We will deduce the lemma from the following more abstract criterion for semisimplicity, which is interesting in its own right – it gives a useful property of semisimple representations, and it is sometimes easier to check because it does not mention irreducible representations.

Lemma 2.2.8 (Semisimplicity criterion). *Let G be a group and let k be a field. A k-representation*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of G is semisimple if and only if, for any subrepresentation $F \subset E$ of ϱ, there exists a complementary subrepresentation, i.e., a G-stable subspace $\tilde{F} \subset E$ such that

$$E = F \oplus \tilde{F}.$$

It is useful to give a name to the second property: one says that a representation ϱ is *completely reducible* if, for any subrepresentation ϱ_1 of ϱ, one can find a complementary one ϱ_2 with

$$\varrho = \varrho_1 \oplus \varrho_2.$$

Proof of Lemma 2.2.7. Let ϱ act on E, and let $F \subset E$ be a subrepresentation. We are going to check that the condition of Lemma 2.2.8 applies to F.[5] Thus let $G \subset F$ be a subrepresentation of F; it is also one of E, hence there exists a subrepresentation $\tilde{G} \subset E$ such that

$$E = G \oplus \tilde{G}.$$

[5] I.e., a subrepresentation of a completely reducible one is itself completely reducible.

Now we claim that $F = G \oplus (F \cap \tilde{G})$, which shows that G has also a stable complement in F – and finishes the proof that F is semisimple. Indeed, G and $(F \cap \tilde{G})$ are certainly in direct sum, and if $v \in F$ and we write $v = v_1 + v_2$ with $v_1 \in G$, $v_2 \in \tilde{G}$, we also obtain

$$v_2 = v - v_1 \in F \cap \tilde{G},$$

because v_1 is also in F. The case of a quotient representation is quite similar and is left to the reader to puzzle... □

Proof of Lemma 2.2.8. Neither direction of the equivalence is quite obvious. We start with a semisimple representation ϱ, acting on E, written as a direct sum

$$E = \bigoplus_{i \in I} E_i$$

of stable subspaces E_i, on which G acts irreducibly, and we consider a stable subspace F. Now we use a standard trick in set-theory: we consider a maximal (for inclusion) subrepresentation $\tilde{F}$ of E such that $F \cap \tilde{F} = 0$, or in other words, such that F and $\tilde{F}$ are in direct sum. Observe that, if the conclusion of the lemma is correct, $\tilde{F}$ must be a full complement of F in E, and we proceed to check this. For every i, consider

$$(F \oplus \tilde{F}) \cap E_i \subset E_i.$$

Since E_i is an irreducible representation of G, this intersection is either 0 or equal to E_i. In fact, it can not be zero, because this would mean that $\tilde{F} + E_i \supsetneq \tilde{F}$ is a larger subrepresentation in direct sum with F, contradicting the definition of $\tilde{F}$. Hence we see that $E_i \subset F \oplus \tilde{F}$ for *all* i, and this means that $F \oplus \tilde{F} = E$.

Now comes the converse: we assume that ϱ, acting on E is non-zero and is completely reducible. We first claim that E contains at least one irreducible subrepresentation: if E has finite dimension, this is Lemma 2.2.5, and otherwise it requires some care but can be done, as explained in Exercise 2.2.10 below.

Now we consider the sum E_1 (not necessarily direct) of all irreducible subrepresentations of E. It is non-zero, as we just observed. In fact, we must have $E_1 = E$, because our assumption implies that $E = E_1 \oplus \tilde{E}_1$ for some other subrepresentation $\tilde{E}_1$, and if $\tilde{E}_1$ were non-zero, it would also contain an irreducible subrepresentation, which contradicts the definition of E_1. Thus E is a *sum* of irreducible subrepresentations, say of E_i, $i \in I$; we proceed to conclude by showing it is a *direct* sum of $(E_i)_{i \in J}$ for some subset $J \subset I$: let J be a maximal subset of I such that the sum of the E_i, $i \in J$, is a direct sum, and let F be the direct sum of those E_i, $i \in J$. For any $i \notin J$, the intersection $E_i \cap F$ can not be zero, as this would allow us to replace J by $J \cap \{i\}$, which is larger than J; hence $E_i \subset F$ for *all* $i \in I$, and hence $E = F$, which is a direct sum of irreducible subrepresentations. □

Remark 2.2.9. In the finite-dimensional case, the last argument can be replaced by an easy induction on $\dim(E)$: if E is not irreducible, we use the assumption to write

$$E \simeq F \oplus F'$$

for some irreducible subspace F and complementary representation F'. The proof of Lemma 2.2.7 really shows that F' is also completely reducible, and since $\dim(F') < \dim(E)$, by induction, we get that F' is also semisimple, and we are done.

Exercise 2.2.10 (Existence of irreducible subrepresentation). Let G be a group, k a field, and $\varrho : G \longrightarrow \mathrm{GL}(E)$ a completely reducible k-representation of G, with $E \neq 0$. We want to show that E contains an irreducible subrepresentation.

(1) Fix a $v \neq 0$. Using Zorn's Lemma, show that there exists a maximal subrepresentation $E_1 \subset E$ (for inclusion) such that $v \notin E_1$.

(2) Write $E = E_1 \oplus E_2$ for some subrepresentation E_2, using the complete reducibility of E. Show that E_2 is irreducible. [Hint: If not, show that $E_2 = E_3 \oplus E_4$ for some non-zero subrepresentations of E_2, and that $v \notin E_1 \oplus E_3$ or $v \notin E_1 \oplus E_4$.]

2.2.4. Tensor product. An equally important construction is the tensor product. Given G and k, and representations ϱ_1 and ϱ_2 of G on k-vector spaces E_1 and E_2, we obtain a representation

$$G \to \mathrm{GL}(E_1 \otimes_k E_2)$$

by sending g to $\varrho_1(g) \otimes \varrho_2(g)$. Thus, by definition, for a pure tensor $v \otimes w \in E_1 \otimes E_2$, we have

$$g \cdot (v \otimes w) = \varrho_1(g)v \otimes \varrho_2(g)w,$$

another pure tensors (but we recall that $E_1 \otimes E_2$ is not simply the space of such pure tensors, but that they generate the tensor product).

The algebraic ("functorial") properties of the tensor operation ensure that this is a group homomorphism. We will denote this representation by $\varrho_1 \otimes \varrho_2$, or sometimes simply by $E_1 \otimes E_2$ when the actions on the vector spaces is clear from context. For the same type of general reasons, all the standard isomorphisms between tensor products such as

$$E_1 \otimes E_2 \simeq E_2 \otimes E_1, \qquad E_1 \otimes (E_2 \otimes E_3) \simeq (E_1 \otimes E_2) \otimes E_3, \qquad E \otimes k \simeq E,$$

are in fact isomorphisms of representations of G, where k in the last equation represents the trivial (one-dimensional) representation of G. In particular, one can define, up to isomorphism, a tensor product involving multiple factors which is independent of the order of the product.

Similarly, we have

$$\varrho \otimes \Bigl(\bigoplus_i \varrho_i\Bigr) \simeq \bigoplus_i (\varrho \otimes \varrho_i).$$

If $\varrho \subset \varrho_1$ is a subrepresentation, tensoring with another representation ϱ_2 gives a subrepresentation

$$\varrho \otimes \varrho_2 \hookrightarrow \varrho_1 \otimes \varrho_2,$$

but one should be careful that, in general, not all subrepresentations of a tensor product are of this form (e.g., because of dimension reasons).

Note finally the relation $\dim(\varrho_1 \otimes \varrho_2) = (\dim \varrho_1)(\dim \varrho_2)$.

2.2.5. Multilinear operations. Besides tensor products, all other multilinear constructions have the "functoriality" property (Proposition 2.2.1) needed to operate at the level of representations of a group. Thus, if $\varrho : G \longrightarrow \mathrm{GL}(E)$ is a k-representation of G, we can construct:

– The symmetric powers $\mathrm{Sym}^m(E)$ of E, for $m \geqslant 0$;

– The alternating powers $\bigwedge^m E$, for $m \geqslant 0$.

In each case, the corresponding operation for endomorphisms of E leads to representations

$$G \longrightarrow \mathrm{GL}(\mathrm{Sym}^m(E)), \qquad G \longrightarrow \mathrm{GL}(\bigwedge^m E),$$

which are called the *m-th symmetric power* and *m-th alternating power* of ϱ, respectively. Taking direct sums leads to representations of G on the symmetric and alternating algebras

$$\mathrm{Sym}(E) = \bigoplus_{m \geqslant 0} \mathrm{Sym}^m(E), \qquad \bigwedge E = \bigoplus_{m \geqslant 0} \bigwedge^m E.$$

From elementary multilinear algebra, we recall that if E is finite-dimensional, the symmetric algebra is infinite-dimensional, but the alternating algebra is not – indeed, $\bigwedge^m E = 0$ if $m > \dim E$. More generally, the dimension of the symmetric and alternating powers is given by

$$\dim \mathrm{Sym}^m(E) = \binom{\dim(E)+m-1}{m}, \qquad \dim \bigwedge^m E = \binom{\dim E}{m}.$$

For instance, if $n = \dim(E)$, we have

$$\dim \mathrm{Sym}^2(E) = \frac{n(n+1)}{2}, \qquad \dim \bigwedge^2 E = \frac{n(n-1)}{2}.$$

Remark 2.2.11 (Symmetric powers as coinvariants). Let E be a k-vector space. For any $m \geqslant 1$, there is a natural representation of the symmetric group $\mathfrak{S}_m$ on the tensor power

$$E^{\otimes m} = E \otimes \cdots \otimes E$$

(with m factors), which is induced by the permutation of the factors, i.e., we have

$$\sigma \cdot (v_1 \otimes \cdots \otimes v_m) = v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(m)}.$$

The classical definition of the m-th symmetric power is

$$\mathrm{Sym}^m(E) = E^{\otimes m}/F$$

where F is the subspace generated by all vectors of the type

$$(v_1 \otimes \cdots \otimes v_m) - (v_{\sigma(1)} \otimes \cdots \otimes v_{\sigma(m)}) = (v_1 \otimes \cdots \otimes v_m) - \sigma \cdot (v_1 \otimes \cdots \otimes v_m)$$

where $v_i \in E$ and $\sigma \in \mathfrak{S}_m$. In other words, in the terminology and notation of Section 2.2.2, we have

$$\mathrm{Sym}^m(E) = E^{\otimes m}_{\mathfrak{S}_m},$$

the space of *coinvariants* of $E^{\otimes m}$ under this action of the symmetric group.

2.2.6. Contragredient, endomorphism spaces. Let ϱ be a k-representation of G, acting on the vector space E. Using the transpose operation, we can then define a representation on the dual space $E' = \mathrm{Hom}_k(E, k)$, which is called the *contragredient* $\tilde{\varrho}$ of ϱ. More precisely, since the transpose reverses products,[6] the contragredient is defined by the rule

$$\langle g \cdot \lambda, v \rangle = \langle \lambda, g^{-1} \cdot v \rangle,$$

for $g \in G$, $\lambda \in E'$ and $v \in E$, using duality-bracket notation, or in other words the linear form $\tilde{\varrho}(g)\lambda$ is the linear form

$$v \mapsto \lambda(\varrho(g^{-1})v).$$

Remark 2.2.12. One way to remember this is to write the definition in the form of the equivalent invariance formula

$$\langle g \cdot \lambda, g \cdot v \rangle = \langle \lambda, v \rangle \tag{2.13}$$

for all $\lambda \in E'$ and $v \in E$.

[6] Equivalently, in the language of Proposition 2.2.1, the assignment $T(E) = E'$ “reverses” arrows in contrast with (2.8), i.e., $T(f \circ g) = T(g) \circ T(f)$, with $T(f)$ the transpose.

We check explicitly that the contragredient is a representation, to see that the inverse (which also reverses products) compensates the effect of the transpose:

$$\langle gh\cdot\lambda, v\rangle = \langle\lambda, (gh)^{-1}\cdot v\rangle = \langle\lambda, h^{-1}g^{-1}\cdot v\rangle = \langle h\cdot\lambda, g^{-1}\cdot v\rangle = \langle g\cdot(h\cdot\lambda), v\rangle$$

for all g, $h\in G$ and $v\in E$.

The following lemma shows how the contragredient interacts with some of the other operations previously discussed:

Lemma 2.2.13. *Let k be a field, G a group.*

(1) *For any k-representations (ϱ_i) of G, we have canonical isomorphisms*

$$\widetilde{\bigoplus_i \varrho_i} \simeq \bigoplus_i \tilde{\varrho}_i.$$

(2) *For any k-representations ϱ_1 and ϱ_2 of G, we have canonical isomorphisms*

$$\widetilde{\varrho_1\otimes\varrho_2} \simeq \tilde{\varrho}_1\otimes\tilde{\varrho}_2.$$

(3) *If a k-representation ϱ of G is such that its contragredient is irreducible, then so is ϱ. Moreover, if ϱ is* finite-dimensional*, then the converse is true, and in fact more generally, if*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

is finite-dimensional, there is an inclusion-reversing bijection between subrepresentations F of ϱ and $\tilde{F}$ of its contragredient, given by

$$F \mapsto F^{\perp} = \{\lambda\in E' \mid \lambda(F) = 0\}$$
$$\tilde{F} \mapsto {}^{\perp}\tilde{F} = \{x\in E \mid \lambda(x) = 0 \textit{ for all } \lambda\in\tilde{F}\}.$$

(4) *For any k-representation ϱ of G of finite dimension, we have a canonical isomorphism $\tilde{\tilde{\varrho}} \simeq \varrho$.*

Proof. Only (3) really requires some words of justification. We observe that both constructions indicated send, in any case, a subrepresentation of ϱ to one of $\tilde{\varrho}$ or conversely: this follows by the definition formula

$$\langle\tilde{\varrho}(g)\lambda, v\rangle = \langle\lambda, \varrho(g^{-1})v\rangle,$$

e.g., if $F\subset E$ is a subrepresentation of ϱ, this implies that $F^{\perp}$ is stable under the contragredient. If ϱ (hence also $\tilde{\varrho}$) is finite-dimensional, standard duality of vector spaces shows that the two operations are inverse to each other. In particular, ϱ is then irreducible if and only if $\tilde{\varrho}$ is.

Without the finite-dimension assumption, we can still argue in this manner to show that if $\tilde{\varrho}$ is irreducible, the original representation is also: for any subrepresentation F of ϱ, the subrepresentation $F^{\perp}$ of $\tilde{\varrho}$ must be either 0 or all of $\tilde{\varrho}$. In the first case, no linear form vanishes on all of F, and that means that F is the whole space; in the second, all linear forms vanish, and this means $F = 0$. Hence ϱ is irreducible. □

Remark 2.2.14. This absence of symmetry in the last parts of this lemma is not surprising because dual spaces of infinite-dimensional vector spaces typically do not behave very well in the absence of topological restrictions.

A well-known isomorphism in linear algebra states that for k-vector spaces E and F, with $\dim(F) < +\infty$, we have

$$\mathrm{Hom}_k(E,F) \simeq E'\otimes F, \tag{2.14}$$

where the isomorphism is induced by mapping a pure tensor $\lambda \otimes v$, with $v \in F$ and $\lambda : E \longrightarrow k$, to the rank 1 homomorphism

$$A_{\lambda,v} : \begin{cases} E & \longrightarrow & E \\ w & \mapsto & \lambda(w)v = \langle \lambda, w \rangle v \end{cases}$$

(because the image of this map lies in the space of finite-rank homomorphisms $E \to F$, we must assume that F has finite dimension to have an isomorphism).

Thus, if

$$\varrho : G \longrightarrow \mathrm{GL}(E), \qquad \tau : G \longrightarrow \mathrm{GL}(F)$$

are k-representations of G, it follows that $\mathrm{Hom}_k(E, F)$ also carries a natural representation of G, defined so that the isomorphism (2.14) is an isomorphism of representations. It is useful to have a more direct description of this action, which is the following: given $A : E \longrightarrow F$ in $\mathrm{Hom}_k(E, F)$ and $g \in G$, define

$$g \cdot A = \tau(g) A \varrho(g)^{-1} : E \longrightarrow F, \tag{2.15}$$

so that the diagram

$$\begin{tikzcd} E \arrow[r, "A"] \arrow[d, "\varrho(g)"'] & F \arrow[d, "\tau(g)"] \\ E \arrow[r, "g\cdot A"] & F \end{tikzcd}$$

commutes.

Concretely, we thus have

$$(g \cdot A)(w) = g \cdot A(g^{-1} \cdot w) \tag{2.16}$$

for all $w \in E$. Note that this definition applies equally to all representations, even if F is not necessarily finite-dimensional. Another way to remember the definition is to write the formula in the form

$$(g \cdot A)(g \cdot w) = g \cdot A(w) \tag{2.17}$$

for $g \in G$ and $w \in E$.

To check that this action is the same as the one described using (2.14), when F has finite dimension, is a simple computation, which (again) the reader should attempt before reading on. Let $\lambda \otimes v$ be a pure tensor in $E' \otimes F$, and $A = A_{\lambda,v}$ the associated homomorphism. Then the rank 1 endomorphism associated to

$$g \cdot (v \otimes \lambda) = g \cdot v \otimes g \cdot \lambda$$

is given by

$$w \mapsto \langle g \cdot \lambda, w \rangle (g \cdot v) = \langle \lambda, g^{-1} w \rangle (g \cdot v) = g \cdot (\langle \lambda, g^{-1} w \rangle v) = g \cdot A(g^{-1} w).$$

These representations on homomorphism spaces are very useful, as we will see in particular in the case of finite groups, as they give a way to "compare" two representations. For instance, note that by (2.17), we have

$$\mathrm{Hom}_k(\varrho_1, \varrho_2)^G = \mathrm{Hom}_G(\varrho_1, \varrho_2) \subset \mathrm{Hom}_k(E, F), \tag{2.18}$$

in other words, the space $\mathrm{Hom}_G(\varrho_1, \varrho_2)$ of G-homomorphisms between ϱ_1 and ϱ_2 is the invariant space in $\mathrm{Hom}_k(\varrho_1, \varrho_2)$. Thus intertwining operators can be detected by looking at invariant linear maps. Part of Schur's Lemma 2.2.4 thus means that $\mathrm{Hom}_k(\pi, \varrho)^G = 0$ if π and ϱ are non-isomorphic irreducible representations.

Remark 2.2.15 (Another action on homomorphism spaces). Given representations ϱ_1 and ϱ_2 of G on E and F, there is another action, say τ, on $\mathrm{Hom}_k(E, F)$ that may come to mind: for $A \in \mathrm{Hom}_k(E, F)$, simply putting

$$(\tau(g)A)(w) = \varrho_2(g)(A(w)), \tag{2.19}$$

for $g \in G$ and $w \in E$, one defines also an action of G on $\mathrm{Hom}_k(E, F)$. This will turn out to be useful below in the proof of Burnside's irreducibility criterion, but it is usually less important than the one previously described. One can guess why: the formula shows that this representation really only involves the representation ϱ_2, and does not "mix" intelligently ϱ_1 and ϱ_2 (a fact that might be obscured from writing the definition in a short-hand like $(g \cdot A)w = g \cdot Aw$; it is also less clear if $\varrho_1 = \varrho_2$, and we consider representations on $\mathrm{End}_k(\varrho_1)$). Concretely, this representation τ is in fact isomorphic to a direct sum of copies of ϱ_2: if $(w_j)_{j \in J}$ is a basis of E, we have an isomorphism

$$\bigoplus_{j \in J} \varrho_2 \longrightarrow \tau \tag{2.20}$$

given by mapping a family (v_j) of vectors in F to the unique linear map such that

$$A(w_j) = v_j.$$

This is a k-linear isomorphism, since (w_j) is a basis, and it is very simple to check that it is an intertwiner.

2.3. Formalism: changing the group

Because composites of homomorphisms are homomorphisms, we see that whenever there exists a group homomorphism

$$H \xrightarrow{\phi} G,$$

it provides a way to associate a k-representation of H to any k-representation

$$\varrho : G \to \mathrm{GL}(E)$$

of G, simply by composition

$$H \xrightarrow{\varrho \circ \phi} \mathrm{GL}(E).$$

The vector space is therefore the same, and the dimension of $\varrho \circ \phi$ is also the same as that of ϱ. Moreover, this operation is compatible with intertwining operators of representations of G (in category-theory language, it is a functor): whenever

$$\Phi : E_1 \longrightarrow E_2$$

is a morphism between k-representations ϱ_1 and ϱ_2 of G on E_1 and E_2 respectively, the linear map Φ is *also* a morphism between $\varrho_1 \circ \phi$ and $\varrho_2 \circ \phi$. Since the morphism of representations of H attached to a composite $\Phi_1 \circ \Phi_2$ is the corresponding composition, one can say that this operation from representations of G to those of H is also *functorial.* In general, this correspondence has no reason to be injective or surjective: some representations of H may not "come from" G in this way, and non-isomorphic representations of G may become isomorphic when "pulled back" to H. The reader is invited to look for (easy!) examples of both phenomena.

When H is a subgroup of G and ϕ is the inclusion, the operation is called, naturally enough, the *restriction* of representations of G to H. Because of this, one uses the standard notation $\mathrm{Res}_H^G(\varrho)$, which we will use even when ϕ is not of this type (note the ambiguity due to the fact that this representation depends on ϕ which is not present in the notation).

EXAMPLE 2.3.1 (Representations of quotients). There is one very common type of "restriction" associated to a non-injective morphism: if $\phi : G \to H$ is surjective, or in other words if $H \simeq G/K$ for some normal subgroup $K \subset G$. One can then describe precisely the representations of G obtained by "restriction" (using ϕ) of those of H:

PROPOSITION 2.3.2 (Representations of a quotient). *Let G be a group and $H = G/K$ a quotient of G, with quotient map*

$$\phi : G \longrightarrow H.$$

For any field k, the map

$$\varrho \mapsto \varrho \circ \phi$$

is a bijection between k-representations ϱ of H and k-representations π of G which are trivial on K, i.e., such that $K \subset \ker(\pi)$.

This is simply a special case of the fact that, for any group Γ, a homomorphism $G \to \Gamma$ factors through K (i.e., is of the form $f \circ \phi$ for some $f : G/K \to \Gamma$) if and only if it is trivial on K.

One of the most basic and important construction of representation theory, and in some sense the first notion that may not be immediately clearly related to notions of linear algebra,[7] is the operation of *induction*, and we will now spend a certain amount of time discussing its basic properties.

This operation proceeds in the direction opposite to restriction: given a homomorphism

$$\phi : H \longrightarrow G,$$

it associates – in a functorial way, i.e., in a way that is natural enough to be compatible with intertwining operators – a k-representation of G to a k-representation of H. When ϕ is the inclusion of a subgroup H of G, this means going from a representation of a subgroup to one of a larger group, which may seem surprising at first. Once more, a reader who has not seen the definition before might want to stop for a few minutes to think if she can come up with a possible way to do this; it is also then useful to read what follows first by assuming ϕ to be an inclusion map.

One defines the induced[8] representation as follows: given

$$\varrho : H \longrightarrow \mathrm{GL}(E),$$

we define first the k-vector space

$$F = \{\varphi : G \to E \mid \varphi(\phi(h)x) = \varrho(h)\varphi(x) \text{ for all } h \in H,\, x \in G\}, \tag{2.21}$$

(which is a vector subspace of the space of functions on G with values in E). In other words, F is the space of E-valued functions on G which happen to transform "like the representation ϱ under H acting on the left". On this vector space F, we now have an action of G, namely the restriction π to F of the regular representation:

$$(\pi(g))\varphi(x) = \varphi(xg)$$

[7] It is, however related to certain tensor products.

[8] It is unfortunate that the terminology "induced" may clash with the use of the adjective "induced" in less formal senses.

for $\varphi \in F$, $g \in G$ and $x \in G$. Indeed, we need only check that F is stable under the regular representation of G; but if $\varphi_1 = \pi(g)\varphi$, we find that

$$\varphi_1(\phi(h)x) = \varphi(\phi(h)xg) = \varrho(h)\varphi(xg) = \varrho(h)\varphi_1(x),$$

for all $h \in H$ and $x \in G$, which means that – as desired – we have $\varphi_1 \in F$ again.

Especially when ϕ is an inclusion, one writes

$$\pi = \operatorname{Ind}_H^G(\varrho)$$

for this induced representation, but as for restriction, we will use it in the general case (keeping in mind the ambiguity that comes from not indicating explicitly ϕ). One may even drop H and G from the notation when they are clear from the context.

Remark 2.3.3. If we take $h \in \ker(\phi)$, the transformation formula in (2.21) for elements of F gives

$$\varphi(x) = \varrho(h)\varphi(x)$$

so that, in fact, any function $\varphi \in F$ takes values in the space $E^{\ker(\phi)}$ of invariants of E under the action of the subgroup $\ker(\phi)$ through ϱ. However, we do not need to state it explicitly in the definition of F, and this avoids complicating the notation. It will reappear in the computation of the dimension of F (Proposition 2.3.8 below). Of course, when ϕ is an inclusion (the most important case), the target space is genuinely E anyway. It is worth observing, however, that as a consequence of Lemma 2.1.9, this subspace $E^{\ker(\phi)}$ is in fact a subrepresentation of E, so that in the condition

$$\varphi(\phi(h)x) = \varrho(h)\varphi(x),$$

the right-hand side also is always in $E^{\ker(\phi)}$.

Example 2.3.4 (Elementary examples of induction). (1) By definition of F, and comparison with the definition of the regular representation, we see that the latter can be expressed as

$$C_k(G) = \operatorname{Ind}_1^G(\mathbf{1}), \tag{2.22}$$

the result of inducing to G the one-dimensional trivial k-representation of the trivial subgroup $1 \to G$.

(2) For further simple orientation, suppose first that $\phi : G \to G$ is the identity. We then have

$$\operatorname{Ind}_G^G(\varrho) \simeq \varrho$$

for any K-representation $\varrho : G \longrightarrow \mathrm{GL}(E)$ of G, the map $F \to E$ giving this isomorphism being simply

$$\varphi \mapsto \varphi(1) \in E,$$

as the reader should make sure to check. (The inverse maps sends $v \in E$ to the function defined by $\varphi(g) = \varrho(g)v$.)

(3) More generally, consider the canonical projection $\phi : G \to G/K$ (the context of Example 2.3.1). For a representation

$$\varrho : G \longrightarrow \mathrm{GL}(E),$$

we then claim that we have

$$\operatorname{Ind}_G^H(\varrho) \simeq E^K$$

with the action of G/K induced by ϱ (note that by Lemma 2.1.9, the subspace E^K is a subrepresentation of E.) This isomorphism is again given by $\varphi \mapsto \varphi(1)$, which – as we have remarked – is a vector in $E^{\ker(\phi)} = E^K$. The reader should again check that this is an isomorphism.

(4) Suppose now that $\phi : G \to G$ is an automorphism. Then, for a representation ϱ of the "source" G, acting on E, the induced representation $\mathrm{Ind}_G^G(\varrho)$ is *not* in general isomorphic to ϱ; rather it is isomorphic to

$$\phi_*\varrho = \varrho \circ \phi^{-1} : G \longrightarrow \mathrm{GL}(E).$$

Indeed, the k-linear isomorphism

$$\Phi \begin{cases} F & \longrightarrow & E \\ \varphi & \mapsto & \varphi(1) \end{cases}$$

satisfies

$$\Phi(\mathrm{reg}(g)\varphi) = \varphi(g) = \varphi(\varrho(\varrho^{-1}(g))) = \varrho(\phi^{-1}(g))\varphi(1) = \phi_*\varrho(\varphi),$$

i.e., it intertwines the induced representation with the representation $\varrho\circ\phi^{-1}$. Incidentally, using again ϕ and seeing ϱ as a representation of the target G, one has of course

$$\mathrm{Res}_G^G(\varrho) = \phi^*\varrho = \varrho \circ \phi.$$

Although this looks like a quick way to produce many "new" representations from one, it is not so efficient in practice because if ϕ is an *inner* automorphism (i.e., if $\phi(g) = xgx^{-1}$ for some fixed $x \in G$), we do have $\phi_*\varrho \simeq \varrho$: by definition, the linear isomorphism $\Phi = \varrho(x)$ satisfies

$$\Phi \circ \phi_*\varrho(g) = \varrho(x)\varrho(x^{-1}gx) = \varrho(g)\Phi$$

for all $g \in G$, and therefore it is an isomorphism $\varphi_*\varrho \longrightarrow \varrho$.

(5) Finally, one can see from the above how to essentially reduce a general induction to one computed using an inclusion homomorphism. Indeed, we have always an isomorphism

$$\mathrm{Ind}_H^G(\varrho) \simeq \mathrm{Ind}_{\mathrm{Im}(\phi)}^G(\phi_*(\varrho^{\ker(\phi)}))$$

where the right-hand side is computed using the inclusion homomorphism $\mathrm{Im}(\phi) \subset G$. This isomorphism is a combination of the previous cases using the factorization

$$H \xrightarrow{\phi_1} H/\ker(\phi) \simeq \mathrm{Im}(\phi) \hookrightarrow G,$$

where the first map is a quotient map, the second the isomorphism induced by ϕ, and the third an injection. (This is also a special case of "induction in steps", see Proposition 2.3.14 below.)

(6) Another important special case of induction occurs when the representation ϱ is one-dimensional, i.e., is a homomorphism

$$H \longrightarrow k^\times.$$

In that case, the space F of $\mathrm{Ind}_H^G(\varrho)$ is a subspace of the space $C_k(G)$ of k-valued functions on G, and since G acts on this space by the regular representation, the induced representation is a subrepresentation of $C_k(G)$, characterized as those functions which transform like ϱ under H:

$$\varphi(\phi(h)x) = \varrho(h)\varphi(x)$$

where now $\varrho(h)$ is just a (non-zero) scalar in k.

This type of example is significant because of the crucial importance of the regular representation. Indeed, it is often a good strategy to (attempt to) determine the irreducible k-representations of a group by trying to find them as being either induced by one-dimensional representations of suitable subgroups, or subrepresentations of such induced representations. We will see this in effect in Chapter 4, in the special case of the groups $\mathrm{GL}_2(\mathbf{F}_q)$, where $\mathbf{F}_q$ is a finite field.

Remark 2.3.5. Although we have given a specific "model" of the induced representation by writing down a concrete vector space on which G acts, one should attempt to think of it in a more abstract way. As we will see in the remainder of the book, many representations constructed differently – or even "given" by nature – turn out to be isomorphic to induced representations, even if the vector space does not look like the one above.

Note also that we have defined induction purely algebraically. As one may expect, in cases where G is an infinite topological group, this definition may require some changes to behave reasonably. The model (2.21) is then a good definition as it can immediately suggest to consider restricted classes of functions on G instead of all of them (see Example 5.2.10.)

The following properties are the most important facts to remember about induction.

Proposition 2.3.6 (Induced representation). *Let k be a field, $\phi : H \to G$ a group homomorphism.*

(1) *For any homomorphism $\varrho_1 \xrightarrow{\Phi} \varrho_2$ of k-representations of H, there is a corresponding homomorphism*

$$\mathrm{Ind}(\Phi) : \mathrm{Ind}_H^G(\varrho_1) \longrightarrow \mathrm{Ind}_H^G(\varrho_2),$$

and this is "functorial": the identity maps to the identity and composites map to composites.

(2) *For any k-representation ϱ_1 of G and ϱ_2 of H, there is a natural isomorphism of k-vector spaces*

$$\mathrm{Hom}_G(\varrho_1, \mathrm{Ind}_H^G(\varrho_2)) \simeq \mathrm{Hom}_H(\mathrm{Res}_H^G(\varrho_1), \varrho_2), \tag{2.23}$$

where we recall that $\mathrm{Hom}_G(\cdot, \cdot)$ denotes the k-vector space of homomorphism between two representations of G.

The last isomorphism is an instance of what is called *Frobenius reciprocity*, and it is an extremely useful result. In fact, in some (precise) sense, this formula characterizes the induced representation, and can be said to more or less define it (see Remark 2.3.15 for an explanation). We will use induction and the Frobenius formula extensively – in particular in the next chapters – to analyze the decomposition of induced representations.

Another remark is that the definition of the induced representation that we chose is the best for handling situations where $[G : H]$ can be infinite. If $[G : H]$ is finite, then another natural (isomorphic) model leads to isomorphisms

$$\mathrm{Hom}_G(\mathrm{Ind}_H^G(\varrho_1), \varrho_2) \simeq \mathrm{Hom}_H(\varrho_1, \mathrm{Res}_H^G(\varrho_2)), \tag{2.24}$$

and those are sometimes considered to be the incarnation of Frobenius reciprocity (see Chapter 4 and, e.g., [**19**, Ch. 5]).

Proof. (1) The induced homomorphism $\Phi_* = \mathrm{Ind}(\Phi)$ is easy to define using the model of the induced representation given above: denoting by F_1, F_2 the spaces corresponding to $\mathrm{Ind}_H^G(\varrho_1)$ and $\mathrm{Ind}_H^G(\varrho_2)$ respectively, we define $\Phi_*(\varphi)$ for $\varphi \in F_1$ to be given by

$$\Phi_*(\varphi)(x) = \Phi(\varphi(x))$$

for $x \in G$. This is a function from G to E_2, by definition, and the relation

$$\Phi_*(\varphi)(\phi(h)x) = \Phi(\varphi(\phi(h)x)) = \Phi(\varrho_1(h)\varphi(x)) = \varrho_2(h)\Phi(\varphi(x))$$

for all $h \in H$ shows that $\Phi_*(\varphi)$ is in the space F_2 of the induced representation of ϱ_2. We leave it to the reader to check that Φ_* is indeed a homomorphism between the representations F_1 and F_2.

(2) Here again there is little that is difficult, except maybe a certain bewildering accumulation of notation, especially parentheses, when checking the details – the reader should however make sure that these checks are done.

Assume that G acts on the space F_1 through ϱ_1, and that H acts on E_2 through ϱ_2. Then the "restriction" of ϱ_1 acts on F_1 through $\varrho_1 \circ \phi$, while the induced representation of ϱ_2 acts on the space F_2 defined as in (2.21).

We will describe how to associate to

$$\Phi : F_1 \longrightarrow F_2,$$

which intertwines ϱ_1 and $\mathrm{Ind}_H^G(\varrho_2)$, a map

$$T(\Phi) : F_1 \longrightarrow E_2$$

intertwining the restriction of ϱ_1 and ϱ_2. We will then describe, conversely, how to start with an intertwiner

$$\Psi : F_1 \longrightarrow E_2$$

and construct another one

$$\tilde{T}(\Psi) : F_1 \longrightarrow F_2,$$

and then it will be seen that $T \circ \tilde{T}$ and $\tilde{T} \circ T$ are the identity morphism, so that T and $\tilde{T}$ give the claimed isomorphisms.

The main point to get from this is that both T and $\tilde{T}$ more or less "write themselves": they express the simplest way (except for putting zeros everywhere!) to move between the desired spaces. One must then check various things (e.g., that functions on G with values in E_2 actually lie in F_2, that the maps are actually intertwiners, that they are reciprocal), but at least once this is done, it is quite easy to recover the definitions.

To begin, given Φ as above and a vector $v \in F_1$, we must define a map $F_1 \longrightarrow E_2$; since $\Phi(v)$ is in F_2, it is a function on G with values in E_2, hence it seems natural to evaluate it somewhere, and the most natural guess is to try to evaluate at the identity element. In other words, we define $T(\Phi)$ to be the map

$$T(\Phi) : \begin{cases} F_1 & \longrightarrow & E_2 \\ v & \mapsto & \Phi(v)(1). \end{cases} \tag{2.25}$$

We can already easily check that $\tilde{\Phi} = T(\Phi)$ is an H-homomorphism (between the restriction of ϱ_1 and ϱ_2): indeed, we have

$$\tilde{\Phi}(h \cdot v) = \tilde{\Phi}(\phi(h)v) = \Phi(\phi(h)v)(1) = \Phi(v)(\phi(h))$$

where the last equality reflects the fact that Φ intertwines ϱ_1 and the induced representation of ϱ_2, the latter acting like the regular representation on F_2. Now because $\Phi(v) \in F_2$, we get

$$\Phi(v)(\phi(h)) = \varrho_2(h)\Phi(v)(1) = \varrho_2(h)\tilde{\Phi}(v)$$

which is what we wanted.

In the other direction, given an H-homomorphism

$$\Psi : F_1 \to E_2,$$

we must construct a map $\tilde{\Psi} = \tilde{T}\Psi$ from F_1 to F_2. Given $v \in F_1$, we need to build a function on G with values in E_2; the function

$$x \mapsto \Psi(\varrho_1(x)v), \tag{2.26}$$

is the most natural that comes to mind, since the values of Ψ are elements of E_2. Thus $\tilde{\Psi}(v)$ is defined to be this function.

We only describe some of the necessary checks required to finish the argument. First, we check that $\varphi = \tilde{\Psi}(v)$ is indeed in F_2: for all $x \in G$ and $h \in H$, we have

$$\varphi(\phi(h)x) = \Psi(\phi(h)\varrho_1(x)v) = \varrho_2(h)\Psi(\varrho_1(x)v) = \varrho_2(h)\varphi(x)$$

(using the fact that Ψ is a homomorphism from $\mathrm{Res}_H^G(\varrho_1)$ to ϱ_2.)

Next, $\tilde{\Psi}$ intertwines ϱ_1 and $\mathrm{Ind}_H^G(\varrho_2)$: for $g \in G$, the function $\tilde{\Psi}(\varrho_1(g)v)$ is

$$x \mapsto \Psi(\varrho_1(xg)v)$$

and this coincides with

$$\mathrm{reg}(g)\tilde{\Psi}(v) = (x \mapsto \tilde{\Psi}(v)(xg)).$$

The remaining property we need is that the two constructions are inverse of each other. If we start with $\Psi \in \mathrm{Hom}_H(F_1, E_2)$, then construct $\tilde{\Psi} = \tilde{T}\Psi$, the definitions (2.25) and (2.26) show that

$$T\tilde{T}\Psi(v) = \tilde{\Psi}(v)(1) = \Psi(v)$$

for all v, i.e., $T \circ \tilde{T}$ is the identity. If we start with $\Phi \in \mathrm{Hom}_G(F_1, F_2)$, define $\Psi = T\Phi$ and $\tilde{\Phi} = \tilde{T}\Psi = \tilde{T}T\Phi$, and unravel the definitions again, we obtain the inescapable conclusion that, given $v \in F_1$, the function $\tilde{\Phi}(v)$ is given by

$$(x \mapsto \Psi(\varrho_1(x)v) = \Phi(\varrho_1(x)v)(1)),$$

and this function of x does coincide with $\Phi(v)$ because

$$\Phi(\varrho_1(x)v) = \mathrm{reg}(x)\Phi(v) = (y \mapsto \Phi(v)(yx)).$$

Thus $\tilde{T} \circ T$ is also the identity, and the proof is finished. □

Example 2.3.7. Let $\varrho_1 = \mathbf{1}$ be the trivial (one-dimensional) representation of G. Then its restriction to H is the trivial representation $\mathbf{1}_H$ of H. By Frobenius reciprocity, we derive

$$\mathrm{Hom}_G(\mathbf{1}_G, \mathrm{Ind}(\varrho_2)) \simeq \mathrm{Hom}_H(\mathbf{1}_H, \varrho_2).$$

Comparing with (2.7), we deduce that there is a (canonical) isomorphism

$$\mathrm{Ind}_H^G(\varrho_2)^G \simeq \varrho_2^H$$

of the invariant subspaces of ϱ_2 and its induced representation.

We now wish to compute the dimension of an induced representation.

Proposition 2.3.8. *Let k be a field, $\phi : H \longrightarrow G$ a group homomorphism. For a k-representation ϱ of H, acting on a space E, we have*

$$\deg(\mathrm{Ind}_H^G(\varrho)) = [G : \mathrm{Im}(\phi)] \dim(E^{\ker(\phi)}).$$

In particular, if H is a subgroup of G, we have

$$\deg(\mathrm{Ind}_H^G(\varrho)) = [G : H] \deg(\varrho).$$

Proof. The idea is very simple: the definition of the space F on which the induced representation acts shows that the value of $\varphi \in F$ at a point x determines the values at all other points of the form $\phi(h)x$, i.e., at all points which are in the same left-coset of G modulo the image of ϕ. Thus there should be $[G : \mathrm{Im}(\phi)]$ independent values of φ; each seems to belong to the space E, but as we have observed in Remark 2.3.3, it is in fact constrained to lie in the possibly smaller space $E^{\ker(\phi)}$.

To check this precisely, we select a set R of representatives of $\mathrm{Im}(\phi)\backslash G$, we let $\tilde{F}$ denote the space of all functions

$$\tilde{\varphi} : R \longrightarrow E^{\ker(\phi)},$$

and we consider the obvious k-linear map

$$F \longrightarrow \tilde{F}$$

defined by restricting functions on G to R (using the remark 2.3.3 to see that this is well-defined). Now we claim that this is an isomorphism of vector spaces, and of course this implies the formula for the dimension of F.

To check the injectivity, we simply observe that if $\varphi \in F$ is identically zero on R, we have

$$\varphi(\phi(h)x) = \varrho(h)\varphi(x) = 0$$

for all $x \in R$ and $h \in H$; since these elements, by definition, cover all of G, we get $\varphi = 0$ (this is really the content of the observation at the beginning of the proof).

For surjectivity, for any $x \in G$, we denote by $r(x)$ the element of R equivalent to x, and we select one $h(x) \in H$ such that

$$x = \phi(h(x))r(x),$$

with $h(x) = 1$ if $x \in R$.

Given an arbitrary function $\tilde{\varphi} : R \to E^{\ker(\phi)}$, we then *define*

$$\varphi(x) = \varphi(\phi(h(x))r(x)) = \varrho(h(x))\tilde{\varphi}(r(x)),$$

which is a well-defined E-valued function on G. Thus φ is equal to $\tilde{\varphi}$ on R; by definition of F, this is in fact the only possible definition for such a function, but we must check that $\varphi \in F$ to conclude. Consider $x \in G$ and $h \in H$; let $y = \phi(h)x$, so that we have the two expressions

$$y = \phi(hh(x))r(x), \qquad y = \phi(h(y))r(y) = \phi(h(y))r(x)$$

since y and x are left-equivalent under $\mathrm{Im}(\phi)$. It follows that $hh(x)$ and $h(y)$ differ by an element (say κ) in $\ker(\phi)$. Thus we get

$$\begin{aligned}\varphi(y) = \varphi(\phi(h(y))r(x)) &= \varrho(h(y))\tilde{\varphi}(r(x)) \\ &= \varrho(\kappa)\varrho(hh(x))\tilde{\varphi}(r(x)) \\ &= \varrho(h)\varrho(h(x))\tilde{\varphi}(r(x))\end{aligned}$$

since κ acts trivially on the space $E^{\ker(\phi)}$, and (as in Lemma 2.1.9) the vector

$$\varrho(hh(x))\tilde{\varphi}(r(x))$$

does belong to it. We are now done because

$$\varphi(\phi(h)x) = \varphi(y) = \varrho(h)\varrho(h(x))\tilde{\varphi}(r(x)) = \varrho(h)\varphi(x)$$

finishes the proof that $\varphi \in F$. □

Remark 2.3.9. From the proof we see that one could have defined the induced representations as the k-vector space of all functions

$$\mathrm{Im}(\phi)\backslash G \longrightarrow E_2^{\ker(\phi)},$$

together with a suitable action of G. However, this "restriction model" of $\mathrm{Ind}_H^G(\varrho)$ is not very convenient because the action of G, by "transport of structure", is not very explicit.

Example 2.3.10 (Minimal index of a proper subgroup). Here is an application of this formula: consider a group G, and a proper subgroup H of G of finite index. We want to obtain a lower bound for its index $[G : H]$. If we fix a field k, then a simple one is given by

$$[G : H] \geqslant \min_{\pi\neq \mathbf{1}} \dim(\pi),$$

where π runs over irreducible, non-trivial, k-representations of G. Indeed, consider the finite-dimensional k-representation

$$\varrho = \mathrm{Ind}_H^G(\mathbf{1}_k),$$

and pick an irreducible subrepresentation π of ϱ. Then we have

$$[G : H] = \dim(\varrho) \geqslant \dim(\pi),$$

as desired.

In many cases, this bound is trivial, but we will see later (see Exercises 4.6.12 and 4.7.3 in Chapter 4) that for the finite groups $\mathrm{SL}_2(\mathbf{F}_p)$, the smallest non-trivial irreducible representation has very large dimension, and there are other important examples.

The degree relation makes it clear, if needed, that the operations of restriction and induction are *not* inverse to each other (as the dimensions of the underlying vector spaces change). In fact, there is no inverse of restriction in general:

Exercise 2.3.11. Show that there is no operation inverse of restriction: there exist subgroups $H \subset G$ and representations of H which are *not* the restriction of any representation of G. [Hint: Even very simple examples will do, and Proposition 2.6.6 below can help.]

Nevertheless, there are relations between restriction and induction, as we have seen with the Frobenius reciprocity formula. Here is another one:

Proposition 2.3.12 (Projection formula). *Let k be a field, and $\phi : H \to G$ a group homomorphism. For a k-representation ϱ_1 of G and a k-representation ϱ_2 of H, we have a natural isomorphism*

$$\mathrm{Ind}_H^G(\varrho_2 \otimes \mathrm{Res}_H^G(\varrho_1)) \simeq \mathrm{Ind}_H^G(\varrho_2) \otimes \varrho_1$$

of representations of G.

As in the case of the Frobenius reciprocity isomorphism (2.23), the proof is not very difficult as the isomorphism can be described explicitly, but full details are a bit tedious. The reader should attempt to guess a homomorphism between the two representations (it is easier to go from right to left here), and after checking that the guess is right, should also try to verify by herself that it satisfies the required properties.[9]

[9] In fact, the details of this and similar proofs are probably not worth trying to read without attempting such a process of self-discovery of the arguments.

Proof. We denote by F_1 the space of ϱ_1, E_2 that of ϱ_2 and F_2 the space (2.21) of the induced representation $\mathrm{Ind}_H^G(\varrho_2)$. Moreover, we denote by τ the representation

$$\tau = \varrho_2 \otimes \mathrm{Res}_H^G(\varrho_1)$$

of H and by $\tilde{F}_2$ the space of

$$\mathrm{Ind}_H^G(\tau) = \mathrm{Ind}_H^G(\varrho_2 \otimes \mathrm{Res}_H^G(\varrho_1)),$$

defined also using (2.21).

The isomorphism of representations of G that is claimed to exist is defined as the k-linear map

$$F_2 \otimes F_1 \xrightarrow{\Phi} \tilde{F}_2$$

induced by

$$\Phi(\varphi \otimes v) = (x \mapsto \varphi(x) \otimes x \cdot v),$$

for $\varphi \in F_2$ and $v \in F_1$, which is indeed a function $G \longrightarrow E_2 \otimes F_1$, the target being the space of τ (in this proof, we write $x \cdot v$ for the action of ϱ_1 on F_1).

It is clear that Φ is well-defined, provided we check that its image does lie in $\tilde{F}_2$. But if $\tilde{\varphi} = \Phi(\varphi \otimes v)$, using the fact that $\varphi \in F_2$, we obtain

$$\begin{aligned}\tilde{\varphi}(\phi(h)x) &= \varphi(\phi(h)x) \otimes (\phi(h)x) \cdot v\\ &= \varrho_2(h)\varphi(x) \otimes \phi(h)(x \cdot v)\\ &= \tau(h)\{\varphi(x) \otimes x \cdot v\}\end{aligned}$$

for all $x \in G$, $h \in H$, which is the property required for a function $G \longrightarrow E_2 \otimes F_1$ to be in $\tilde{F}_2$.

We will now check that Φ is a G-isomorphism. First, the fact that it is a homomorphism is straightforward, as it can be checked on the generating tensors $\varphi \otimes v$. Let $\tilde{\varphi} = \Phi(\varphi \otimes v)$ and $g \in G$; then we have

$$(g \cdot \tilde{\varphi})(x) = \tilde{\varphi}(xg) = \varphi(xg) \otimes (xg) \cdot v$$

which we can also write as

$$\varphi_1(x) \otimes x \cdot w$$

where $\varphi_1(x) = \varphi(xg) = g \cdot \varphi(x)$ and $w = g \cdot v$, or in other words as

$$\Phi(\varphi_1 \otimes w)(x) = \Phi(g \cdot (\varphi \otimes v))(x),$$

as desired.

It remains, to conclude, to prove that Φ is a k-linear isomorphism. Here a little trick is needed, since pure tensors are not enough. We fix a basis (v_j) of F_1 (it could be infinite, of course). Then, for any $x \in G$, a vector w of $E_2 \otimes F_1$ can be written *uniquely* as a linear combination

$$w = \sum_j w_j(x) \otimes (x \cdot v_j) \tag{2.27}$$

for some $w_j(x) \in E_2$. This is simply because, for every x, the vectors $(x \cdot v_j)_j$ also form a basis of F_1.

We now first show the injectivity of Φ: any element of $F_2 \otimes F_1$ can be expressed as

$$\sum_j \varphi_j \otimes v_j$$

for some functions $\varphi_j \in F_2$. Let us assume such an element is in $\ker(\Phi)$. This means that for all $x \in G$, we have

$$\sum_j \varphi_j(x) \otimes (x \cdot v_j) = 0 \in E_2 \otimes F_1.$$

Thus by the uniqueness of the representations (2.27), we get

$$\varphi_j(x) = 0$$

for all j, or in other words $\varphi_j = 0$ for all j, so that $\ker(\Phi) = 0$.

We now come to surjectivity. Let $\tilde{\varphi} \in \tilde{F}_2$ be given. Again by the observation above, for any $x \in G$, we can write uniquely[10]

$$\tilde{\varphi}(x) = \sum_j \tilde{\varphi}_j(x) \otimes (x \cdot v_j),$$

thus defining coefficient functions $\tilde{\varphi}_j : G \to E_2$. We now will show that – because $\tilde{\varphi} \in \tilde{F}_2$ – each $\tilde{\varphi}_j$ is in fact in F_2, which will ensure that

$$\tilde{\varphi} = \sum_j \Phi(\tilde{\varphi}_j \otimes v_j)$$

is in the image of Φ, which is therefore surjective.

The condition $\tilde{\varphi} \in \tilde{F}_2$ means that

$$\tilde{\varphi}(\phi(h)x) = \tau(h)\tilde{\varphi}(x)$$

for all $h \in H$ and $x \in G$. The left-hand side is

$$\sum_j \tilde{\varphi}_j(\phi(h)x) \otimes (\phi(h)x \cdot v_j)$$

by definition, while the right-hand side is

$$\begin{aligned}(\varrho_2 \otimes \operatorname{Res} \varrho_1)(h)\tilde{\varphi}(x) &= \sum_j \varrho_2(h)\tilde{\varphi}_j(x) \otimes \{\phi(h) \cdot (x \cdot v_j)\} \\ &= \sum_j \varrho_2(h)\tilde{\varphi}_j(x) \otimes (\phi(h)x \cdot v_j).\end{aligned}$$

Comparing using the uniqueness of (2.27), with x replaced by $\phi(h)x$, we find that, for all j, we have

$$\tilde{\varphi}_j(\phi(h)x) = \varrho_2(h)\tilde{\varphi}_j(x)$$

and this does state that each coefficient function $\tilde{\varphi}_j$ is in F_2. □

Remark 2.3.13. If ϕ is an injective homomorphism and the groups G and H are finite, then all spaces involved are finite-dimensional. Since Proposition 2.3.8 shows that both sides of the projection formula are of degree $[G : H] \deg(\varrho_1) \deg(\varrho_2)$, the injectivity of Φ is sufficient to finish the proof.

Yet another property of induction (and restriction), which is quite important, is the following:

[10] This is the trick: using (2.27) for a varying x, not for a single fixed basis.

Proposition 2.3.14 (Transitivity). *Let k be a field and let*

$$H_2 \xrightarrow{\phi_2} H_1 \xrightarrow{\phi_1} G$$

be group homomorphisms, and let $\phi = \phi_1 \circ \phi_2$. For any k-representations ϱ_2 of H_2 and ϱ of G, we have canonical isomorphisms

$$\mathrm{Res}^{H_1}_{H_2}(\mathrm{Res}^{G}_{H_1} \varrho) \simeq \mathrm{Res}^{G}_{H_2}(\varrho), \qquad \mathrm{Ind}^{G}_{H_1}(\mathrm{Ind}^{H_1}_{H_2} \varrho_2) \simeq \mathrm{Ind}^{G}_{H_2}(\varrho_2).$$

Proof. As far as the restriction is concerned, this is immediate from the definition. For induction, the argument is pretty much of the same kind as the ones we used before: defining maps both ways is quite simple and hard to miss, and then one merely needs to make various checks to make sure that everything works out; we will simplify those by omitting the homomorphisms in the notation.

So here we go again: let E, F_1, F_2, F denote, respectively, the spaces of the representations

$$\varrho_2, \quad \mathrm{Ind}^{H_1}_{H_2} \varrho_2, \quad \mathrm{Ind}^{G}_{H_1}(\mathrm{Ind}^{H_1}_{H_2} \varrho_2), \quad \mathrm{Ind}^{G}_{H_2}(\varrho_2),$$

so that we must define a G-isomorphism

$$T : F \longrightarrow F_2.$$

Note that F is a space of functions from G to E, and F_2 a space of functions from G to F_1. We define T as follows: given $\varphi \in F$, a function from G to E, it is natural to consider

$$\mathrm{reg}(g)\varphi = (x \mapsto \varphi(xg)),$$

the image of φ under the regular representation on E-valued functions. But $T(\varphi)$ must be F_1-valued to be in F_2, and F_1 is a space of functions from H_1 to E. Hence we define

$$T(\varphi)(g) = (\mathrm{reg}(g)\varphi) \circ \phi_1,$$

the "restriction" to H_1 of this function on G.

We can then check that $T(\varphi)$ is, in fact, F_1-valued; if we assume that the group homomorphisms involved are inclusions of subgroups, this amounts to letting $\eta = T(\varphi)(g)$ and writing

$$\eta(h_2h_1) = \mathrm{reg}(g)\varphi(h_2h_1) = \varphi(h_2h_1g) = \varrho(h_2)\varphi(h_1g) = \varrho(h_2)\eta(h_1),$$

for $h_i \in H_i$, using of course in the middle the assumption that φ is in F (again, the confusing mass of parentheses is unlikely to make much sense until the reader has tried and succeeded independently to do it).

Now we should check that $T(\varphi)$ is not only F_1-valued, but also lies in F_2, i.e., transforms under H_1 like the induced representation $\mathrm{Ind}^{H_1}_{H_2}(\varrho)$. We leave this to the reader: this is much helped by the fact that the action of H_1 on this induced representation is also the regular representation.

Next, we must check that T is an intertwining operator; but again, both F and F_2 carry actions which are variants of the regular representation, and this should not be surprising – we therefore omit it...

The final step is the construction of the inverse $\tilde{T}$ of T.[11] We now start with $\psi \in F_2$ and must define a function from G to E; unraveling in two steps, we set

$$\tilde{T}(\psi)(g) = \psi(g)(1)$$

[11] If the vector spaces are finite dimensional and the homomorphisms are inclusions, note that it is quite easy to check that T is injective, and since the dimensions of F and F_2 are both $[G : H_2] \dim \varrho$, this last step can be shortened.

($\psi(g)$ is an element of F_1, i.e., a function from H_1 to E, and we evaluate that at the unit of H_1...) Taking $g \in G$ and $h_2 \in H_2$, denoting $\varphi = \tilde{T}(\psi)$, we again let the reader check that the following

$$\varphi(h_2 g) = \psi(h_2 g)(1) = (\mathrm{reg}(h_2)\psi(g))(1) = \psi(g)(h_2) = \varrho(h_2)\psi(g)(1) = \varrho(h_2)\varphi(g),$$

makes sense, and means that $\tilde{T}(\psi)$ is in F.

Now we see that $\tilde{T}T(\varphi)$ is the function which maps $g \in G$ to

$$(\mathrm{reg}(g)\varphi)(1) = \varphi(g),$$

in other words $\tilde{T} \circ T$ is the identity. Rather more abstrusely, if $\psi \in F_2$, $\varphi = \tilde{T}(\psi)$ and $\tilde{\psi} = T(\varphi)$, we find for $g \in G$ and $h_1 \in H_1$ that

$$\begin{aligned}\tilde{\psi}(g)(h_1) &= (\mathrm{reg}(g)\varphi)(h_1) = \varphi(h_1 g)\\ &= \psi(h_1 g)(1) = (\mathrm{reg}(h_1)\psi(g))(1)\\ &= \psi(g)(h_1)\end{aligned}$$

(where we use the fact that, on F_2, H_1 acts through the regular representation), which indicates that $T \circ \tilde{T}$ is also the identity. Thus both are reciprocal isomorphisms. □

Remark 2.3.15 (Functoriality saves time). At this point, conscientious readers may well have become bored and annoyed at this “death of a thousand checks”. And there are indeed at least two ways to avoid much (if not all) of the computations we have done. One will be explained in Section 3.1, and we sketch the other now, since the reader is presumably well motivated to hear about abstract nonsense if it cuts on the calculations.

The keyword is the adjective “natural” (or “canonical”) that we attributed to the isomorphisms (2.23) of Frobenius Reciprocity. In one sense, this is intuitive enough: the linear isomorphism

$$\mathrm{Hom}_G(\varrho_1, \mathrm{Ind}_H^G(\varrho_2)) \longrightarrow \mathrm{Hom}_H(\mathrm{Res}_H^G(\varrho_1), \varrho_2),$$

defined in the proof of Proposition 2.3.6 certainly feels natural. But we now take this more seriously, and try to give rigorous sense to this sentence.

The point is the following fact: a representation ϱ of G is determined, up to isomorphism, by the data of all homomorphism spaces

$$V(\pi) = \mathrm{Hom}_G(\pi, \varrho)$$

where π runs over k-representations of G, *together* with the data of the maps

$$V(\pi) \xrightarrow{V(\Phi)} V(\pi')$$

associated to any (reversed!) G-homomorphism $\pi' \xrightarrow{\Phi} \pi$ by mapping

$$(\Psi\,:\,\pi \to \varrho) \in V(\pi)$$

to

$$V(\Phi)(\Psi) = \Psi \circ \Phi.$$

To be precise:

Fact. Suppose that ϱ_1 and ϱ_2 are k-representations of G, and that for any representation π, there is given a k-linear isomorphism

$$I(\pi)\,:\,\mathrm{Hom}_G(\pi, \varrho_1) \longrightarrow \mathrm{Hom}_G(\pi, \varrho_2),$$

in such a way that all diagrams

$$\begin{tikzcd}
\mathrm{Hom}_G(\pi, \varrho_1) \arrow[r] \arrow[d] & \mathrm{Hom}_G(\pi, \varrho_2) \arrow[d] \\
\mathrm{Hom}_G(\pi', \varrho_1) \arrow[r] & \mathrm{Hom}_G(\pi', \varrho_2)
\end{tikzcd}$$

commute for any $\Phi : \pi' \to \pi$, where the vertical arrows, as above, are given by $\Psi \mapsto \Psi \circ \Phi$. Then ϱ_1 and ϱ_2 are isomorphic, and in fact there exists a unique isomorphism

$$\varrho_1 \xrightarrow{I} \varrho_2$$

such that $I(\pi)$ is given by $\Psi \mapsto I \circ \Psi$ for all π.

Let us first see why this is useful. When dealing with induction, the point is that it tells us that an induced representation $\mathrm{Ind}_H^G(\varrho)$ is characterized, up to isomorphism, by the Frobenius Reciprocity isomorphisms (2.23). Indeed, the latter tells us, simply from the data of ϱ, what any G-homomorphism space

$$\mathrm{Hom}_G(\pi, \mathrm{Ind}_H^G(\varrho))$$

is supposed to be. And the fact above says that there can be only one representation with a "given" homomorphism groups $\mathrm{Hom}_G(\pi, \varrho')$.

More precisely, the behavior under morphisms must be compatible – the fact above gives:

Fact. Let $\phi : H \to G$ be a group-homomorphism and let ϱ be a k-representation of H. There exists, up to isomorphism of representations of G, *at most* one k-representation ϱ' of G with k-linear isomorphisms

$$i(\pi) : \mathrm{Hom}_G(\pi, \varrho') \longrightarrow \mathrm{Hom}_H(\mathrm{Res}(\pi), \varrho)$$

such that the diagrams

$$\begin{tikzcd}
\mathrm{Hom}_G(\pi, \varrho') \arrow[r, "i(\pi)"] \arrow[d] & \mathrm{Hom}_H(\mathrm{Res}(\pi), \varrho) \arrow[d] \\
\mathrm{Hom}_G(\pi', \varrho') \arrow[r, "i(\pi')"] & \mathrm{Hom}_H(\mathrm{Res}(\pi'), \varrho)
\end{tikzcd}$$

commute for $\pi' \xrightarrow{\Phi} \pi$ a G-homomorphism, where the vertical arrows are again $\Psi \mapsto \Psi \circ \Phi$, on the left, and $\Psi \mapsto \Psi \circ \mathrm{Res}(\Phi)$ on the right (restriction of Φ to H)

Readers are invited to check that the (explicit) isomorphisms

$$i(\pi) : \mathrm{Hom}_G(\pi, \mathrm{Ind}_H^G(\varrho)) \longrightarrow \mathrm{Hom}_H(\mathrm{Res}_H^G(\pi), \varrho),$$

that we constructed (based on the explicit model (2.21)) are such that the diagrams

$$\begin{tikzcd}
\mathrm{Hom}_G(\pi, \mathrm{Ind}_H^G(\varrho)) \arrow[r, "i(\pi)"] \arrow[d] & \mathrm{Hom}_H(\mathrm{Res}_H^G(\pi), \varrho) \arrow[d] \\
\mathrm{Hom}_G(\pi', \mathrm{Ind}_H^G(\varrho)) \arrow[r, "i(\pi')"] & \mathrm{Hom}_H(\mathrm{Res}_H^G(\pi'), \varrho)
\end{tikzcd} \tag{2.28}$$

commute (these are the same as the ones above, with $\varrho' = \mathrm{Ind}(\varrho)$). This is the real content of the observation that those are "natural". Thus the *construction* of (2.21) proved the existence of the induced representation defined by "abstract" Frobenius reciprocity...

Now we can see that the transitivity of induction is just a reflection of the – clearly valid – transitivity of restriction. Consider

$$H_2 \xrightarrow{\phi_2} H_1 \xrightarrow{\phi_1} G$$

as in the transitivity formula, and consider a representation ϱ of H_2, as well as

$$\varrho_1 = \mathrm{Ind}_{H_1}^{G}(\mathrm{Ind}_{H_2}^{H_1}(\varrho)), \qquad \varrho_2 = \mathrm{Ind}_{H_2}^{G}(\varrho).$$

According to Frobenius reciprocity applied twice or once, respectively, we have, for all representations π of G, k-linear isomorphisms

$$\mathrm{Hom}_G(\pi, \varrho_1) \simeq \mathrm{Hom}_{H_1}(\mathrm{Res}_{H_1}^{G}(\pi), \mathrm{Ind}_{H_2}^{H_1}(\varrho)) \simeq \mathrm{Hom}_{H_2}(\mathrm{Res}_{H_2}^{H_1}(\mathrm{Res}_{H_1}^{G}(\pi)), \varrho)$$

and

$$\mathrm{Hom}_G(\pi, \varrho_2) \simeq \mathrm{Hom}_{H_2}(\mathrm{Res}_{H_2}^{G}(\pi), \varrho),$$

hence by comparison and the "obvious" transitivity of restriction, we obtain isomorphisms

$$I(\pi) \,:\, \mathrm{Hom}_G(\pi, \varrho_1) \simeq \mathrm{Hom}_G(\pi, \varrho_2).$$

The reader should easily convince herself (and then check!) that these isomorphisms satisfy the compatibility required in the claim to deduce that ϱ_1 and ϱ_2 are isomorphic – indeed, this is a "composition" or "tiling" of the corresponding facts for the diagrams (2.28).

At first sight, this may not seem much simpler than what we did earlier, but a second look reveals that we did not use *anything* relating to k-representations of G except the existence of morphisms, the identity maps and the composition operations! In particular, there is no need whatsoever to know an explicit model for the induced representation.

Now we prove the claim: take $\pi = \varrho_1$; then $\mathrm{Hom}_G(\pi, \varrho_1) = \mathrm{Hom}_G(\varrho_1, \varrho_1)$. We may not know much about the general existence of homomorphisms, but certainly this space contains the identity of ϱ_1. Hence we obtain an element

$$I = I(\varrho_1)(\mathrm{Id}_{\varrho_1}) \in \mathrm{Hom}_G(\varrho_1, \varrho_2).$$

Then – this looks like a cheat – this I *is* the desired isomorphism! To see this – but first try it! –, we check first that $I(\pi)$ is given, as claimed, by pre-composition with I for any π. Indeed, $I(\pi)$ is an isomorphism

$$\mathrm{Hom}_G(\pi, \varrho_1) \longrightarrow \mathrm{Hom}_G(\pi, \varrho_2).$$

Take an element $\Phi \,:\, \pi \to \varrho_1$; we can then build the associated commutative square

$$\begin{tikzcd} \mathrm{Hom}_G(\varrho_1, \varrho_1) \arrow[r] \arrow[d] & \mathrm{Hom}_G(\varrho_1, \varrho_2) \arrow[d] \\ \mathrm{Hom}_G(\pi, \varrho_1) \arrow[r] & \mathrm{Hom}_G(\pi, \varrho_2) \end{tikzcd} .$$

Take the element Id_{ϱ_1} in the top-left corner. If we follow the right-then-down route, we get, by definition the element

$$I(\pi)(\mathrm{Id}) \circ \Phi = I \circ \Phi \in \mathrm{Hom}_G(\pi, \varrho_2).$$

But if we follow the down-then-right route, we get $I(\pi)(\mathrm{Id} \circ \Phi)$, and hence the commutativity of these diagrams says that, for all Φ, we have

$$I(\pi)(\Phi) = I \circ \Phi, \tag{2.29}$$

which is what we had claimed.

We now check that I is, indeed, an isomorphism, by exhibiting an inverse. The construction we used strongly suggests that

$$J = I(\varrho_2)^{-1}(\mathrm{Id}_{\varrho_2}) \in \mathrm{Hom}_G(\varrho_2, \varrho_1),$$

should be what we need (where we use that $I(\varrho_2)$ is an isomorphism, by assumption). Indeed, tautologically, we have

$$I(\varrho_2)(J) = \mathrm{Id}_{\varrho_2},$$

which translates, from the formula (2.29) we have just seen, to

$$I \circ J = \mathrm{Id}_{\varrho_2}.$$

Now we simply exchange the role of ϱ_1 and ϱ_2 and replace $I(\pi)$ by its inverse; then I and J are exchanged, and we get also

$$J \circ I = \mathrm{Id}_{\varrho_1}.$$

Why did we not start with this "functorial" language? Partly this is a matter of personal taste and partly of wanting to show very concretely what happens – especially if the reader does (or has done...) all computations on her own, par of the spirit of the game will have seeped in. Moreover, in some of the more down-to-earth applications of these games with induction and its variants, it may be quite important to know what the "canonical maps" actually are. The functorial language does usually give a way to compute them, but it may be more direct to have written them down as directly as we did.

To conclude with the general properties of induction, we leave the proofs of the following lemma to the reader:

Lemma 2.3.16. *Let k be a field, let $\phi : H \longrightarrow G$ be a group homomorphism. For any representations ϱ and ϱ_i of H, we have natural isomorphisms*

$$\widetilde{\mathrm{Ind}_H^G(\varrho)} \simeq \mathrm{Ind}_H^G(\tilde{\varrho}),$$

and

$$\mathrm{Ind}_H^G\Bigl(\bigoplus_{i\in I} \varrho_i\Bigr) \simeq \bigoplus_{i\in I} \mathrm{Ind}_H^G(\varrho_i).$$

The corresponding statements for the restriction are also valid, and equally easy to check. On the other hand, although the isomorphism

$$\mathrm{Res}_H^G(\varrho_1 \otimes \varrho_2) \simeq \mathrm{Res}_H^G(\varrho_1) \otimes \mathrm{Res}_H^G(\varrho_2),$$

is immediate, it is usually definitely false (say when ϕ is injective but is not an isomorphism) that

$$\mathrm{Ind}_H^G(\varrho_1 \otimes \varrho_2), \quad \mathrm{Ind}_H^G(\varrho_1) \otimes \mathrm{Ind}_H^G(\varrho_2),$$

are isomorphic, for instance because the degrees do not match (from left to right, they are given by

$$[G : H] \deg(\varrho_1) \deg(\varrho_2), \qquad [G : H]^2 (\deg \varrho_1)(\deg(\varrho_2)$$

respectively).

We conclude this longish section with another type of "change of groups". Fix a field k and two groups G_1 and G_2. Given k-representations ϱ_1 and ϱ_2 of G_1 and G_2, acting on E_1 and E_2 respectively, we can define a representation of the direct product $G_1 \times G_2$ on the tensor product $E_1 \otimes E_2$: for pure tensors $v_1 \otimes v_2$ in $E_1 \otimes E_2$, we let

$$(\varrho_1 \boxtimes \varrho_2)(g_1, g_2)(v_1 \otimes v_2) = \varrho_1(g_1)v_1 \otimes \varrho_2(g_2)v_2,$$

which extends by linearity to the desired action, sometimes called the *external tensor product*

$$\varrho_1 \boxtimes \varrho_2 : G_1 \times G_2 \longrightarrow \mathrm{GL}(E_1 \otimes E_2).$$

Of course, the dimension of this representation is again $(\dim \varrho_1)(\dim \varrho_2)$. In particular, it is clear that not all representations of $G_1 \times G_2$ can be of this type, simply because their dimensions might not factor non-trivially. However, in some cases, *irreducible* representations must be external tensor products of irreducible representations of the factors.

PROPOSITION 2.3.17 (Irreducible representations of direct products). *Let k be a field, and let G_1, G_2 be two groups. If ϱ is a finite-dimensional irreducible k-representation of $G = G_1 \times G_2$, then there exist irreducible k-representations ϱ_1 of G_1 and ϱ_2 of G_2, respectively, such that*

$$\varrho \simeq \varrho_1 \boxtimes \varrho_2 \ ;$$

moreover, ϱ_1 and ϱ_2 are unique, up to isomorphism of representations of their respective groups.

Conversely, if ϱ_1 and ϱ_2 are irreducible finite-dimensional k-representations of G_1 and G_2, respectively, the external tensor product $\varrho_1 \boxtimes \varrho_2$ is an irreducible representation of $G_1 \times G_2$.

The proof of this requires some preliminary results, so we defer it to Section 2.7 (Proposition 2.3.17 and Exercise 2.7.29).

REMARK 2.3.18 (Relation with the ordinary tensor product). Consider a group G; there is a homomorphism (which is injective)

$$\phi \begin{cases} G & \longrightarrow & G \times G \\ g & \mapsto & (g, g). \end{cases}$$

If ϱ_1 and ϱ_2 are k-representations of G, the definitions show that $\mathrm{Res}_G^{G\times G}(\varrho_1 \boxtimes \varrho_2) = \varrho_1 \otimes \varrho_2$.

2.4. Formalism: changing the field

We will not say much about changing the field. Clearly, whenever K is an extension of k, we can turn a k-representation

$$G \longrightarrow \mathrm{GL}(E)$$

into a representation over K, by composing with the group homomorphism

$$\mathrm{GL}(E) \longrightarrow \mathrm{GL}(E \otimes_k K)$$

which, concretely (see the next section also), can be interpreted simply by saying that a matrix with coefficients in the subfield k of K can be seen as a matrix with coefficients in K, i.e., by looking at the inclusion

$$\mathrm{GL}_n(k) \hookrightarrow \mathrm{GL}_n(K).$$

If ϱ is a representation of G over a field K, and it is *isomorphic* to a representation arising in this manner from a k-representation, for some subfield k of K, one customarily says that ϱ can be *defined over k*.

A certain property $\mathcal{P}$ of ϱ, such as semisimplicity, irreducibility, existence of non-trivial direct sum decompositions, may not exist over the small field k, but may become possible over an extension K, for instance an algebraic closure $\bar{k}$ of k. In this last case, one says that ϱ "has $\mathcal{P}$ absolutely", or "has $\mathcal{P}$ geometrically".

Example 2.4.1. Consider the (infinite) abelian group $G = \mathbf{R}/2\pi\mathbf{Z}$, and the 2-dimensional real representation given by

$$\varrho : \begin{cases} G \longrightarrow \mathrm{GL}_2(\mathbf{R}) \\ \theta \mapsto \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \end{cases},$$

which corresponds to the action of $\mathbf{R}/2\pi\mathbf{Z}$ on the real plane by rotation of a given angle. This makes it clear that this is a homomorphism, as trigonometric identities can also be used to check. Also, this interpretation makes it clear that ϱ is irreducible: there is no non-zero (real) subspace of $\mathbf{R}^2$ which is stable under ϱ, except $\mathbf{R}^2$ itself.

However, this irreducibility breaks down when extending the base field to $\mathbf{C}$: indeed, on $\mathbf{C}^2$, we have

$$\varrho(\theta)\begin{pmatrix} 1 \\ i \end{pmatrix} = \begin{pmatrix} \cos\theta + i\sin\theta \\ -\sin\theta + i\cos\theta \end{pmatrix} = (\cos\theta + i\sin\theta)\begin{pmatrix} 1 \\ i \end{pmatrix},$$

and

$$\varrho(\theta)\begin{pmatrix} 1 \\ -i \end{pmatrix} = \begin{pmatrix} \cos\theta - i\sin\theta \\ \sin\theta - i\cos\theta \end{pmatrix} = (\cos\theta - i\sin\theta)\begin{pmatrix} 1 \\ -i \end{pmatrix},$$

so that $\mathbf{C}^2$, under the action of G through ϱ, splits as a direct sum

$$\mathbf{C}^2 = \begin{pmatrix} 1 \\ i \end{pmatrix}\mathbf{C} \oplus \begin{pmatrix} 1 \\ -i \end{pmatrix}\mathbf{C}$$

of two complex lines which are both subrepresentations, one of them isomorphic to the one-dimensional complex representation

$$\begin{cases} G \to \mathrm{GL}(\mathbf{C}) \simeq \mathbf{C}^\times \\ \theta \mapsto e^{i\theta} \end{cases}$$

and the other to

$$\begin{cases} G \to \mathbf{C}^\times \\ \theta \mapsto e^{-i\theta} \end{cases}$$

(its conjugate, in a fairly obvious sense). Hence, one can say that ϱ is *not absolutely irreducible.*

Another way to change the field, which may be more confusing, is to apply *automorphisms* of k. Formally, this is not different: we have an automorphism $\sigma : k \longrightarrow k$, and we compose ϱ with the resulting homomorphism

$$\varrho_\sigma : G \xrightarrow{\varrho} \mathrm{GL}(E) \longrightarrow \mathrm{GL}(E \otimes_k k),$$

where we have to be careful to see k, in the second argument of the tensor product, as given with the k-algebra structure σ. Concretely, $E_\sigma = E \otimes_k k$ is the k-vector space with the same underlying abelian group E, but with scalar multiplication given by

$$\alpha \cdot v = \sigma(\alpha)v \in E.$$

Here again matrix representations may help understand what happens: a basis (v_i) of E is still a basis of E_σ but, for any $g \in G$, the matrix representing $\varrho(g)$ in the basis (v_i) of E_σ is obtained by applying σ^{-1} to all coefficients of the matrix that represents $\varrho(g)$. Indeed, for any i, we have

$$\varrho(g)v_i = \sum_j \alpha_j v_j = \sum_j \sigma^{-1}(\alpha_j) \cdot v_j$$

so that the (j,i)-th coefficient of the matrix for $\varrho(g)$ is α_j, while it is $\sigma^{-1}(\alpha_j)$ for $\varrho_\sigma(g)$.

This operation on representations can be interesting because ϱ and ϱ_σ are usually *not* isomorphic as representations, despite the fact that they are closely related. In particular, there is a bijection between the subrepresentations of E and those of E_σ (given by $F \mapsto F_\sigma$), and hence ϱ and ϱ_σ are simultaneously irreducible or not irreducible, semisimple or not semisimple.

Example 2.4.2 (Complex conjugate). Consider $k = \mathbf{C}$. Although $\mathbf{C}$, considered as an abstract field, has many automorphisms, the only continuous ones, and therefore the most important, are the identity and the complex conjugation $\sigma : z \mapsto \bar{z}$. It follows therefore that any time we have a complex representation $G \longrightarrow \mathrm{GL}(E)$, where E is a $\mathbf{C}$-vector space, there is a naturally associated "conjugate" representation $\bar{\varrho}$ obtained by applying the construction above to the complex conjugation. From the basic theory of characters (Corollary 2.7.32 below), one can see that $\bar{\varrho}$ is isomorphic to ϱ if and only if the function $g \mapsto \operatorname{Tr} \varrho(g)$ is real-valued. This can already be checked when ϱ is one-dimensional, since $\bar{\varrho}$ is then the conjugate function $G \to \mathbf{C}$, which equals ϱ if and only if ϱ is real-valued. In particular, the examples in (2.4), (2.5) or (2.6) lead to many cases of representations where ϱ and $\bar{\varrho}$ are not isomorphic.

Field extensions are the only morphisms for fields. However, there are sometimes other possibilities to change fields, which are more subtle.

2.5. Matrix representations

We have emphasized in Definition 2.1.1 the abstract view where a representation is seen as a linear action of G on a k-vector space E. However, in practice, if one wishes to *compute* with representations, one will select a fixed basis of E and express ϱ as a homomorphism

$$\varrho^{\boldsymbol{m}} : G \longrightarrow \mathrm{GL}_n(k), \qquad n = \dim(E),$$

that maps g to the matrix representing $\varrho(g)$ in the chosen basis. Indeed, this is how we did in the cases of the Example in (2.12) and in Example 2.4.1.

Although such matrix representations can be awkward when used exclusively, it is useful and important to know how to express in these terms the various operations on representations that we have described previously. These concrete descriptions may also help clarify these operations, especially for readers less familiar with abstract algebra. We will explain this here fairly quickly.

For a direct sum $\varrho_1 \oplus \varrho_2$, we concatenate bases $(e_1, \ldots, e_n)$ of E_1 and $(f_1, \ldots, f_m)$ of E_2 to obtain a basis

$$(e_1, \ldots, e_n, f_1, \ldots, f_m)$$

in which the representation $\varrho_1 \oplus \varrho_2$ takes the form of block-diagonal matrices

$$g \mapsto \left(\begin{array}{c|c} \varrho_1^{\boldsymbol{m}}(g) & 0 \\ \hline 0 & \varrho_2^{\boldsymbol{m}}(g) \end{array}\right)$$

of size $m+n$. Corresponding to a short exact sequence

$$0 \to E_1 \longrightarrow E \xrightarrow{\Phi} E_2 \to 0$$

of representations, which does not necessarily split, we select a basis $(e_1, \ldots, e_n)$ of the subspace E_1 of E, and we extend it to a basis $(e_1, \ldots, e_n, f_1, \ldots, f_m)$, $m = \dim(E_2)$, of E. Then

$$(f_1', \ldots, f_m') = (\Phi(f_1), \ldots, \Phi(f_m))$$

is a basis of E_2 and we get in these bases a block-triangular matrix representation of ϱ acting on E:

$$g \mapsto \left(\begin{array}{c|c} \varrho_1^{\boldsymbol{m}}(g) & \star \\ \hline 0 & \varrho_2^{\boldsymbol{m}}(g) \end{array}\right) \tag{2.30}$$

(where $\varrho_1^{\boldsymbol{m}}$ is the matrix-representation in $(e_1, \ldots, e_n)$ and $\varrho_2^{\boldsymbol{m}}$ the one in $(f_1', \ldots, f_m')$).

In the case of a tensor product $\varrho = \varrho_1 \otimes \varrho_2$, one usually represents it in the basis of pure tensors $\delta_{i,j} = e_i \otimes f_j$. If it is ordered as follows

$$(\delta_{1,1}, \delta_{1,2}, \ldots, \delta_{2,1}, \ldots, \delta_{n,m}),$$

and we denote by $A = (a_{i,j})_{1 \leqslant i,j \leqslant n}$ the matrix $\varrho_1^{\boldsymbol{m}}(g)$ and by B the matrix $\varrho_2^{\boldsymbol{m}}(g) = (b_{i,j})$, then $\varrho^{\boldsymbol{m}}(g)$ is a block matrix with n rows and columns, and square blocks of size m given by

$$\left(\begin{array}{c|c|c|c} a_{1,1}B & a_{1,2}B & \ldots & a_{1,n}B \\ \hline \vdots & \vdots & & \vdots \\ \hline a_{n,1}B & \ldots & \ldots & a_{n,n}B \end{array}\right).$$

The matrix-representation of the contragredient of a representation ϱ is also easy to describe: we have

$$\tilde{\varrho}^{\boldsymbol{m}}(g) = {}^t\varrho^{\boldsymbol{m}}(g)^{-1},$$

the inverse-transpose homomorphism.

The case of the restriction to a subgroup is immediate: the matrices of the restriction do not change. For induction, the situation is more involved, and we will only give examples in the next chapters.

2.6. Examples

We collect here some more examples of representations.

2.6.1. Binary forms and invariants. Let k be any field. For an integer $m \geqslant 0$, we denote by V_m the vector space of polynomials in $k[X, Y]$ which are homogeneous of degree m, i.e., the k-subspace of $k[X, Y]$ generated by the monomials

$$X^i Y^{m-i}, \qquad 0 \leqslant i \leqslant m.$$

In fact, these monomials are independent, and therefore form of basis of V_m. In particular $\dim V_m = m + 1$.

If we take $G = \mathrm{SL}_2(k)$, we can let G act on V_m by

$$(\varrho_m(g)f)(X, Y) = f((X, Y) \cdot g),$$

where $(X, Y) \cdot g$ denotes the right-multiplication of matrices; in other words, if

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

we have

$$(g \cdot f)(X, Y) = f(aX + cY, bX + dY)$$

(one just says that G acts on V_m by linear change of variables).

The following theorem will be proved partly in Example 2.7.9):

Theorem 2.6.1 (Irreducible representations of SL_2). *For $k = \mathbf{C}$, the representations ϱ_m, for $m \geqslant 0$, are irreducible representation of* $\mathrm{SL}_2(\mathbf{C})$. *In fact, ϱ_m is then an irreducible representation of the subgroup* $\mathrm{SU}_2(\mathbf{C}) \subset \mathrm{SL}_2(\mathbf{C})$.

On the other hand, if k is a field of non-zero characteristic p, the representation ϱ_p is not irreducible.

We only explain the last statement here: if k has characteristic p, consider the subspace $W \subset V_p$ spanned by the monomials X^p and Y^p. Then $W \neq V_p$ (since $\dim V_p = p+1$), and V_p is a subrepresentation. Indeed, we have

$$(g \cdot X^p) = (aX + cY)^p = a^p X^p + c^p Y^p \in W, \quad (g \cdot Y^p) = (aX + cY)^p = b^p X^p + d^p Y^p,$$

by the usual properties of the p-th power operation in characteristic p (i.e., the fact that the binomial coefficients $\binom{p}{j}$ are divisible by p for $1 \leqslant j \leqslant p-1$). One can also show that $W \subset V_p$ does not have a stable complementary subspace, so that V_p is not semisimple in characteristic p.

We now consider only the case $k = \mathbf{C}$. It is elementary that ϱ_m is isomorphic to the m-th symmetric power of ϱ_1 for all $m \geqslant 0$. Hence we see here a case where, using multilinear operations, all irreducible (finite-dimensional) representations of a group are obtained from a "fundamental" one. We also see here an elementary example of a group which has irreducible finite-dimensional representations of arbitrarily large dimension. (In fact, $\mathrm{SL}_2(\mathbf{C})$ also has many infinite-dimensional representations which are irreducible, in the sense of representations of topological groups explained in Section 3.2 in the next chapter.)

Exercise 2.6.2 (Matrix representation). (1) Compute the matrix representation for ϱ_2 and ϱ_3, in the bases (X^2, XY, Y^2) and (X^3, X^2Y, XY^2, Y^3) of V_2 and V_3, respectively.

(2) Compute the kernel of ϱ_2 and ϱ_3.

A very nice property of these representations – which turns out to be crucial in Quantum Mechanics – illustrates another important type of results in representation theory:

Theorem 2.6.3 (Clebsch-Gordan formula). *For any integers $m \geqslant n \geqslant 0$, the tensor product $\varrho_m \otimes \varrho_n$ is semisimple and decomposes as*

$$\varrho_m \otimes \varrho_n \simeq \varrho_{m+n} \oplus \varrho_{m+n-2} \oplus \cdots \oplus \varrho_{m-n}. \tag{2.31}$$

One point of this formula is to illustrate that, if one knows some irreducible representations of a group, one may well hope to be able to construct or identify others by trying to decompose the tensor products of these representations into irreducible components (if possible); here, supposing one knew only the "obvious" representations $\varrho_0 = \mathbf{1}$ and ϱ_1 (which is just the inclusion $\mathrm{SL}_2(\mathbf{C}) \longrightarrow \mathrm{GL}_2(\mathbf{C})$), we see that all other representations ϱ_m arise by taking tensor products iteratively and decomposing them, e.g.,

$$\varrho_1 \otimes \varrho_1 = \varrho_2 \oplus \mathbf{1}, \qquad \varrho_2 \otimes \varrho_1 = \varrho_3 \oplus \varrho_1, \quad \text{etc.}$$

Proof. Both sides of the Clebsch-Gordan formula are trivial when $m = 0$. Using induction on m, we then see that it is enough to prove that

$$\varrho_m \otimes \varrho_n \simeq \varrho_{m+n} \oplus (\varrho_{m-1} \otimes \varrho_{n-1}) \tag{2.32}$$

for $m \geqslant n \geqslant 1$.

At least a subrepresentation isomorphic to $\varrho_{m-1} \otimes \varrho_{n-1}$ is not too difficult to find. Indeed, first of all, the tensor product $\varrho_m \otimes \varrho_n$ can be interpreted concretely by a representation on the space $V_{m,n}$ of polynomials in four variables X_1, Y_1, X_2, Y_2 which are homogeneous of degree m with respect to (X_1, Y_1), and of degree n with respect to the other variables, where the group $\mathrm{SL}_2(\mathbf{C})$ acts by simultaneous linear change of variable on the two sets of variables, i.e.,

$$(g \cdot f)(X_1, Y_1, X_2, Y_2) = f\left((X_1, Y_1)g, (X_2, Y_2)g\right)$$

for $f \in V_{m,n}$. This G-isomorphism

$$V_m \otimes V_n \longrightarrow V_{m,n}$$

is induced by

$$(X^i Y^{m-i}) \otimes (X^j Y^{n-j}) \mapsto X_1^i Y_1^{m-i} X_2^j Y_2^{n-j}$$

for the standard basis vectors.

Using this description, we have a linear map

$$\Delta \quad \begin{cases} V_{m-1,n-1} & \longrightarrow \quad V_{m,n} \\ \quad f & \mapsto \quad (X_1Y_2 - X_2Y_1)f \end{cases}$$

which is a G-homomorphism: if we view the factor $X_1Y_2 - X_2Y_1$ has a determinant

$$\delta = \begin{vmatrix} X_1 & X_2 \\ Y_1 & Y_2 \end{vmatrix},$$

it follows that

$$\delta((X_1, Y_1)g, (X_2, Y_2)g) = \delta(X_1, X_2, Y_1, Y_2)\det(g) = \delta(X_1, X_2, Y_1, Y_2)$$

for $g \in \mathrm{SL}_2(\mathbf{C})$. Moreover, it should be intuitively obvious that Δ is injective, but we check this rigorously: if $f \neq 0$, it has degree $d \geqslant 1$ with respect to some variable, say X_1, and then X_1Y_2f has degree $d+1$ with respect to X_1, while X_2Y_1f remains of degree d, and therefore $X_1Y_2f \neq X_2Y_1f$.

Now we describe a stable complement to the image of Δ. To justify a bit the solution, note that $\mathrm{Im}(\Delta)$ only contains polynomials f such that $f(X, Y, X, Y) = 0$. Those for which this property fails must be recovered. We do this by defining W to be the representation generated by the single vector

$$e = X_1^m X_2^n,$$

i.e., the linear span in $V_{m,n}$ of the translates $g \cdot e$. To check that it has the required property, we look on W at the linear map "evaluating f when both sets of variables are equal" suggested by the remark above. It is given by

$$T \quad \begin{cases} W \longrightarrow V_{m+n} \\ f \mapsto f(X, Y, X, Y) \end{cases}$$

(since a polynomial of the type $f(X, Y, X, Y)$ with $f \in V_{m,n}$ is homogeneous of degree $m+n$), and we notice that it is an intertwiner with ϱ_{m+n}. Since e maps to X^{m+n} which is non-zero, and ϱ_{m+n} is irreducible (Theorem 2.6.1; although the proof of this will be given only later, the reader will have no problem checking that there is no circularity), Schur's Lemma 2.2.4 proves that T is surjective.

We now examine W more closely. Writing $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, we have

$$g \cdot e = (aX_1 + bY_1)^m (aX_2 + bY_2)^n = \sum_{0 \leqslant j \leqslant m+n} a^j b^{m+n-j} \varphi_j(X_1, Y_1, X_2, Y_2)$$

for some $\varphi_j \in V_{m,n}$. We deduce that the space W, spanned by the vectors $g \cdot e$, is contained in the span of the φ_j, and hence that $\dim W \leqslant m+n+1$. But since $\dim \varrho_{m+n} = m+n+1$, we must have equality, and in particular T is an isomorphism.

Since $\dim V_{m-1,n-1} + \dim W = mn+m+n+1 = \dim V_{m,n}$, there only remains to check that $V_{m-1,n-1} \oplus W = V_{m,n}$ to conclude that (2.32) holds. But the intersection $V_{m-1,n-1} \cap W$ is zero, since $f(X,Y,X,Y) = 0$ for $f \in V_{m-1,n-1}$, while $f(X,Y,X,Y) = Tf \neq 0$ for a non-zero $f \in W$... □

In Corollary 5.6.4 in Chapter 5, we will see that the Clebsch-Gordan formula for the subgroup $\mathrm{SU}_2(\mathbf{C})$ (i.e., seeing each ϱ_m as restricted to $\mathrm{SU}_2(\mathbf{C})$) can be proved – at least at the level of *existence* of an isomorphism! – in a few lines using character theory. However, the proof above has the advantage that it "explains" the decomposition, and can be used to describe concretely the subspaces of $V_m \otimes V_n$ corresponding to the subrepresentations of $\varrho_m \otimes \varrho_n$.

Now, in a slightly different direction, during the late 19th and early 20th Century, a great amount of work was done on the topic called *invariant theory*, which in the (important) case of the invariants of $\mathrm{SL}_2(\mathbf{C})$ can be described as follows: one considers, for some $m \geqslant 0$, the algebra $S(V_m)$ of all polynomial functions on V_m; the group G acts on $S(V_m)$ according to

$$(g \cdot \phi)(f) = \phi(\varrho_m(g^{-1})f).$$

and hence $S(V_m)$ is also a representation of G (it is infinite-dimensional, but splits as a direct sum of the homogeneous components of degree $d \geqslant 0$, which are finite-dimensional). Then one tries to understand the subalgebra $S(V_m)^G$ of all G-invariant functions on V_m, in particular, to understand the (finite) dimensions of the homogeneous pieces $S(V_m)^G_d$ of invariant functions of degree d.

For instance, if $m = 2$, so that V_2 is the space of binary quadratic forms, one can write any $f \in V_2$ as

$$f = a_0X^2 + 2a_1XY + a_2Y^2,$$

and then $S(V_2) \simeq \mathbf{C}[a_0, a_1, a_2]$ is the polynomial algebra in these coordinates. One invariant springs to mind: the *discriminant*

$$\Delta(a_0, a_1, a_2) = a_1^2 - a_0a_2$$

of a binary quadratic form. One can then show that $S(V_2)^G \simeq \mathbf{C}[\Delta]$ is a polynomial algebra in the discriminant. For $m = 3$, with

$$f = a_0X^3 + 3a_1X^2Y + 3a_2XY^2 + a_3Y^3,$$

one can prove that $S(V_3)^G \simeq \mathbf{C}[\Delta_3]$, where

$$\Delta_3 = a_0^2a_3^2 - 6a_0a_1a_2a_3 + 4a_0a_2^3 - 3a_1^2a_2^2 + 4a_1^3a_3.$$

The search for explicit descriptions of the invariant spaces $S(V_m)^G$ – and similar questions for other linear actions of groups like $\mathrm{SL}_m(\mathbf{C})$ on homogeneous polynomials in more variables – was one of main topics of the classical theory of invariants, which was extremely popular during the 19-th century (see, e.g., [**39**, Ch. 3] for a modern presentation). These questions are in fact very hard if one wishes to give concrete answers: currently, explicit generators of $S(V_m)^G$ (as an algebra) seem to be known only for $m \leqslant 10$. For $m = 9$, one needs 92 invariants to generate $S(V_m)^G$ as an algebra (see [**6**]; these generators are *not* algebraically independent).

2.6.2. Permutation representations. At the origin of group theory, a group G was often seen as a "permutation group", or in other words, as a subgroup of the group $\mathfrak{S}_X$ of all bijections of some set X (often finite). Indeed, any group G can be identified with a subgroup of $\mathfrak{S}_G$ by mapping $g \in G$ to the permutation $h \mapsto gh$ of the underlying set G (i.e., mapping g to the g-th row of the "multiplication table" of the group law on G). More generally, one may consider any action of G on a set X, i.e., any homomorphism

$$\begin{cases} G \longrightarrow \mathfrak{S}_X \\ g \mapsto (x \mapsto g \cdot x) \end{cases}$$

as a "permutation group" analogue of a linear representation. Such actions, even if X is not a vector space, are often very useful means of investigating the properties of a group. There is always an associated linear representation which encapsulates the action by "linearizing it": given any field k, denote by E_X the k-vector space generated by basis vectors e_x indexed by the elements of the set X, and define

$$\varrho : G \longrightarrow \mathrm{GL}(E_X)$$

by linearity using the rule

$$\varrho(g)e_x = e_{g\cdot x}$$

which exploits the action of G on X. Since $g \cdot (h \cdot x) = (gh) \cdot x$ (the crucial defining condition for an action!), we see that ϱ is, indeed, a representation of G. It has dimension $\dim \varrho = |X|$, by construction.

Example 2.6.4. (1) The representation (denoted ϱ_1) in (2.3) of G on the space $k(G)$ spanned by G is simply the permutation representation associated to the left-action of G on itself by multiplication.

(2) If $H \subset G$ is a subgroup of G, with finite index, and $X = G/H$ is the finite set of right cosets of G modulo H, with the action given by

$$g \cdot (xH) = gxH \in G/H,$$

the corresponding permutation representation ϱ is isomorphic to the induced representation

$$\mathrm{Ind}_H^G(\mathbf{1}).$$

Indeed, the space for this induced representation is given by

$$F = \{\varphi \,:\, \varphi(hg) = \varphi(g) \text{ for all } h \in H\},$$

with the left-regular representation. This space has a basis given by the functions φ_x which are the characteristic functions of the *left* cosets Hx. Moreover

$$\mathrm{reg}(g)\varphi_x = \varphi_{xg^{-1}}$$

(the left-hand side is non-zero at those y where $yg \in Hx$, i.e., $y \in Hxg^{-1}$), which means that mapping

$$\varphi_x \mapsto e_{x^{-1}}$$

gives a linear isomorphism $F \longrightarrow E_X$, which is now an intertwiner.

A feature of all permutation representations is that if X is finite then, unless $|X| = 1$, they are never irreducible: the element

$$\sum_{x \in X} e_x \in E_X$$

is in an invariant vector.

Exercise 2.6.5. If ϱ is the permutation representation associated to the action on G/H, for $H \subset G$ of finite index, show that ϱ^G is spanned by this invariant vector, and explain how to recover it as the image of an explicit element

$$\Phi \in \mathrm{Hom}_G(\mathbf{1}, \varrho)$$

constructed using Frobenius reciprocity.

2.6.3. Generators and relations. From an abstract point of view, one may try to describe representations of a group G by writing down a presentation of G, i.e., a set $\underline{g} \subset G$ of generators, together with the set $\underline{r}$ of all relations between the elements of $\underline{g}$, relations being seen as (finite) words involving the $g \in \underline{g}$ – a situation which one summarizes by writing

$$G \simeq \langle \underline{g} \mid \underline{r} \rangle.$$

Then one can see that for a given field k and dimension $d \geqslant 1$, it is equivalent to give a d-dimensional (matrix) representation

$$G \longrightarrow \mathrm{GL}_d(k)$$

or to give a family

$$(x_g)_{g \in \underline{g}}$$

of invertible matrices in $\mathrm{GL}_d(k)$, such that "all relations in $\underline{r}$ hold", i.e., if a given $r \in \underline{r}$ is given by a word

$$r = g_1 \cdots g_\ell$$

(with g_i in the free group generated by $\underline{g}$), we should ensure that

$$x_{g_1} \cdots x_{g_\ell} = 1$$

in the matrix group $\mathrm{GL}_d(k)$.

This description is usually not very useful for practical purposes if the group G is given, since it is often the case that there is no particularly natural choice of generators and relation to use. However, it does have some purpose.

For instance, if we restrict to representations of dimension $d = 1$, since $\mathrm{GL}_1(k) = k^\times$ is abelian (and no group $\mathrm{GL}_d(k)$ is, for any k, when $d \geqslant 2$), and since for any group G and abelian group A, there is a canonical bijection between homomorphisms $G \rightarrow A$ and homomorphisms $G/[G, G] \rightarrow A$, we derive:

Proposition 2.6.6 (1-dimensional representations). *Let G be a group, and let $G^{ab} = G/[G, G]$ be the* abelianization *of G. For any field k, the 1-dimensional representations of G correspond with the homomorphisms*

$$G^{ab} \longrightarrow k^\times.$$

In particular, if G is perfect, *i.e., if $[G, G] = G$, any non-trivial representation of G, over any field k, has dimension at least 2.*

The last part of this proposition applies in many cases. For instance, if $d \geqslant 2$ and k is any field, $\mathrm{SL}_d(k)$ is known to be perfect except when $d = 2$ and $k = \mathbf{F}_2$ or $k = \mathbf{F}_3$. Thus no such group has a non-trivial one-dimensional representation.

One can also make use of this approach to provide more examples of groups with "a lot" of representations. Indeed, if G is a group where there are no relations at all between a set $\underline{g}$ of generators (i.e., a free group), it is equivalent to give a homomorphism $G \longrightarrow \mathrm{GL}(E)$ as to give elements x_g in $\mathrm{GL}(E)$ indexed by the generators $g \in \underline{g}$. Moreover, two such representations given by $x_g \in \mathrm{GL}(E)$ and $y_g \in \mathrm{GL}(F)$ are isomorphic if and only if these elements are (globally) conjugate, i.e., if there exists a linear isomorphism $\Phi : E \to F$ such that

$$x_g = \Phi^{-1} y_g \Phi$$

for all $g \in \underline{g}$.

Here is a slight variant that makes this very concrete. Consider the group $G = \mathrm{PSL}_2(\mathbf{Z})$ of matrices of size 2 with integral coefficients and determinant 1, modulo the subgroup $\{\pm 1\}$. Then G is not free, but it is known to be generated by the (image modulo $\{\pm 1\}$ of the) two elements

$$g_1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad g_2 = \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$$

in such a way that the only relations between the generators are

$$g_1^2 = 1, \qquad g_2^3 = 1$$

(i.e., G is a *free product* of $\mathbf{Z}/2\mathbf{Z}$ and $\mathbf{Z}/3\mathbf{Z}$).

Hence it is equivalent to give a representation $\mathrm{PSL}_2(\mathbf{Z}) \longrightarrow \mathrm{GL}(E)$ or to give two elements x, $y \in \mathrm{GL}(E)$ such that $x^2 = 1$ and $y^3 = 1$.

Yet another use of generators and relations is in showing that there exist groups for which certain representations *do not exist*: in that case, it is enough to find some abstract presentation where the relations are incompatible with matrix groups. Here is a concrete example:

Theorem 2.6.7 (Higman–Baumslag; "non-linear finitely generated groups exist"). *Let G be the group with 2 generators a, b subject to the relation*

$$a^{-1} b^2 a = b^3.$$

Then, whatever the field k, there exists no faithful linear representation

$$G \longrightarrow \mathrm{GL}(k)$$

where E is a finite-dimensional k-vector space.

The first example of such a group was constructed by Higman; the example here is due to Baumslag (see [**29**]), and is an example of a family of groups called the Baumslag-Solitar groups which have similar presentations with the exponents 2 and 3 replaced by arbitrary integers.

We will only give a sketch, dependent on some fairly deep facts of group theory.

Sketch of proof. We appeal to the following two results:

– (Malcev's Theorem) If k is a field and $G \subset \mathrm{GL}_d(k)$ is a finitely generated group, then for any $g \in G$, there exists a finite quotient $G \longrightarrow G/H$ such that g is non-trivial modulo H.

– (The "Identitätssatz" of Magnus, or Britton's Lemma; see, e.g., [**33**, Th. 11.81]) Let G be a finitely presented group with a *single relation* (a one-relator group); then one can decide algorithmically if a word in the generators represents or not the identity element of G.

Now we are going to check that G fails to satisfy the conclusion of Malcev's Theorem, and therefore has no finite-dimensional representation over any field.

To begin with, iterating the single relation leads to

$$a^{-k} b^{2^k} a^k = b^{3^k}$$

for all $k \geqslant 1$. Now assume $G \xrightarrow{\pi} G/H$ is a finite quotient of G, and let $\alpha = \pi(a)$, $\beta = \pi(b)$. Taking k to be the order of α in the finite group G/H, we see that

$$\beta^{2^k - 3^k} = 1,$$

i.e., the order of β divides $2^k - 3^k$. In particular, this order is coprime with 2, and this implies that the map $\gamma \mapsto \gamma^2$ is surjective on the finite cyclic group generated by β. Thus β is a power of β^2. Similarly, after conjugation by a, the element $b_1 = a^{-1}ba$ is such that $\beta_1 = \pi(b_1)$ is a power of β_1^2.

But now we observe that $\beta_1^2 = \pi(a^{-1}b^2a) = \pi(b^3) = \beta^3$. Hence β_1 is a power of β^3, and in particular it commutes with β, so that

$$\pi([b_1, b]) = \beta_1\beta\beta_1^{-1}\beta^{-1} = 1$$

and this is valid in any finite quotient.

Now Britton's Lemma [**33**, Th. 11.81] implies that the word

$$c = [b_1, b] = b_1bb_1^{-1}b^{-1} = a^{-1}baba^{-1}b^{-1}ab^{-1}$$

is non-trivial in G.[12] Thus $c \in G$ is an element which is non-trivial, but becomes so in any finite quotient of G. This is the desired conclusion. □

Remark 2.6.8. Concerning Malcev's Theorem, the example to keep in mind is the following: a group like $\mathrm{SL}_d(\mathbf{Z}) \subset \mathrm{GL}_d(\mathbf{C})$ is finitely generated and one can check that it satisfies the desired condition simply by using the reduction maps

$$\mathrm{SL}_d(\mathbf{Z}) \longrightarrow \mathrm{SL}_d(\mathbf{Z}/p\mathbf{Z})$$

modulo primes. Indeed, for any fixed $g \in \mathrm{SL}_d(\mathbf{Z})$, if $g \neq 1$, we can find some prime p larger than the absolute values of all coefficients of g, and then g is certainly non-trivial modulo p. The proof of Malcev's Theorem is based on similar ideas (though of course one has to use more general rings than $\mathbf{Z}$).

Note that if one does not insist for finitely-generated counterexamples, it is easier to find non-linear groups – for instance, "sufficiently big" abelian groups will work.

2.7. Some general results

In this section, we will prove some of the basic facts about representations. Some of them will, for the first time, require some restrictive assumption, namely either that we consider finite-dimensional representations, or that the base field k be algebraically closed.

2.7.1. The Jordan-Hölder-Noether theorem.

We first discuss a generalization of the classical Jordan-Hölder theorem of group theory, which explains in which sense irreducible representations are in fact "building blocks" of all representations, at least for the finite-dimensional case.

Theorem 2.7.1 (Jordan-Hölder-Noether theorem). *Let G be a group, k a field, and*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

a k-representation of G.

[12] In the language explained in Rotman's book, G is an HNN extension for $A = 2\mathbf{Z}$, $B = 3\mathbf{Z}$, isomorphic subgroups of $\mathbf{Z} = \langle b \rangle$, with stable letter a; thus the expression for c contains no "pinch" $a^{-1}b^2a$ or ab^3a^{-1} as a subword, and Britton's Lemma deduces from this that $c \neq 1$.

(1) *If $E \neq 0$ and E is finite-dimensional, there exists a finite increasing sequence of subrepresentations*

$$0 = E_0 \subset E_1 \subset E_2 \subset \cdots \subset E_{k-1} \subset E_k = E$$

of E such that, for all i, $1 \leqslant i \leqslant k$, the quotient representations E_i/E_{i-1} are irreducible. Such sequences are called composition series*, and the E_i/E_{i+1} are called the* composition factors*.*

(2) *If ϱ admits any finite composition series,*[13] *then any two such sequences are equivalent, in the following sense: the number of terms are the same, and the irreducible representations which appear – i.e., the composition factors – are isomorphic, up to a permutation. In other words, for any sequence (E_i) as above, and for any irreducible k-representation π of G, the integer*

$$n_\pi(\varrho) = |\{i \mid E_i/E_{i-1} \simeq \pi\}|$$

is the same.

The uniqueness part of the statement may be considered, to some extent, as analogue of the fundamental theorem of arithmetic: a factorization of an integer into prime powers is unique, but only up to permutation of the primes.

Remark 2.7.2. (1) The result is often simply called the Jordan-Hölder Theorem, but according to H. Weyl [**44**], the extension to representations is due to E. Noether.

(2) By definition, any composition factor of a representation ϱ is isomorphic to a quotient ϱ_1/ϱ_2 where $\varrho_2 \subset \varrho_1 \subset \varrho$ are subrepresentations of ϱ. More generally, any representation of this type, not necessarily irreducible, is called a *subquotient* of ϱ.

Proof. The existence part (1) is easy, by dimension arguments: since $E \neq 0$, we can select an irreducible subrepresentation E_1 (a subrepresentation of minimal non-zero dimension), then – if $E_1 \neq E$ – a subrepresentation $E_2 \supsetneq E_1$ of minimal dimension, etc. For dimension reasons, each quotient E_i/E_{i-1} is then irreducible, and the process terminates in finitely many steps because $\dim(E) < +\infty$.

The uniqueness is more important. Assume that we have two sequences

$$0 = E_0 \subset E_1 \subset E_2 \subset \cdots \subset E_{k-1} \subset E_k = E,$$
$$0 = F_0 \subset F_1 \subset F_2 \subset \cdots \subset F_{\ell-1} \subset F_\ell = E$$

with irreducible quotients F_j/F_{j-1} and E_i/E_{i-1}. We proceed to use the second one to insert (apparent) steps between the successive subspaces of the first sequence – and conversely. Precisely, for $0 \leqslant i \leqslant k-1$, let

$$E_{i,j} = E_i + (E_{i+1} \cap F_j), \qquad 0 \leqslant j \leqslant \ell$$

and for $0 \leqslant j \leqslant \ell - 1$, let

$$F_{j,i} = F_j + (F_{j+1} \cap E_i), \qquad 0 \leqslant i \leqslant k.$$

Then we have, e.g,

$$E_i = E_{i,0} \subset E_{i,1} \subset \cdots \subset E_{i,\ell-1} \subset E_{i,\ell} = E_{i+1},$$

and a similar refinement between F_j and F_{j+1}.

By construction each $E_{i,j}$ and $F_{i,j}$ is a subrepresentation of E. Now, for each i, $0 \leqslant i \leqslant k-1$, observe that since there is no proper intermediate subrepresentations

[13] This may happen even if $\dim(E)$ is infinite!

between E_i and E_{i+1} (this would contradict the fact that E_{i+1}/E_i is irreducible), there exists a *unique* index j, $0 \leqslant j \leqslant \ell - 1$, for which

$$E_{i,j} = E_i, \qquad E_{i,j+1} = E_{i+1},$$

hence with

$$E_{i+1}/E_i = (E_i + (E_{i+1} \cap F_{j+1}))/(E_i + (E_{i+1} \cap F_j)).$$

There is a certain symmetry in this between i and j; in fact, by a standard isomorphism theorem, there is a canonical isomorphism

$$(2.33)\quad (E_i + (E_{i+1} \cap F_{j+1}))/(E_i + (E_{i+1} \cap F_j)) \simeq (F_j + (E_{i+1} \cap F_{j+1}))/(F_j + (E_i \cap F_{j+1})).$$

(see below for a reminder on this).

The right-hand side is none other than $F_{j,i+1}/F_{j,i}$. The latter is therefore non-zero, but for the same reason as before, there is a single step of the interpolated sequence between F_j and F_{j+1} that can be non-zero, and it must satisfy

$$F_{j,i+1}/F_{j,i} \simeq F_{j+1}/F_j.$$

In other words, for each successive irreducible quotient of the *first* sequence, we have associated – canonically, as it turns out – a well-defined irreducible quotient of the second sequence. This gives a map

$$\begin{cases} \{1, \ldots, k\} & \longrightarrow & \{1, \ldots, \ell\} \\ \quad i & \mapsto & j \end{cases}$$

which is injective, because j is characterized by the isomorphism

$$F_{j+1}/F_j \simeq F_{j,i+1}/F_{j,i},$$

which holds for a single index i.

Reversing the role of the two sequences, we obtain the equality $k = \ell$, and then a natural bijection between the irreducible quotients in the first sequence and those of the second. □

Remark 2.7.3 (About the "standard" isomorphism). The isomorphism (2.33) can be expressed as

$$(E + (\tilde{E} \cap \tilde{F}))/(E + (\tilde{E} \cap F)) \simeq (F + (\tilde{E} \cap \tilde{F}))/(F + (E \cap \tilde{F}))$$

for subrepresentations $E \subset \tilde{E}$, $F \subset \tilde{F}$ of some ambient space. This is induced simply by quotienting the reciprocal linear maps

$$e + \tilde{g} \mapsto \tilde{g}$$
$$f + \tilde{g} \mapsto \tilde{g}$$

for $e \in E$, $f \in F$ and $\tilde{g} \in \tilde{E} \cap \tilde{F}$. Indeed, formally at least, these maps are inverse to each other: losing e from left to right is no problem because the E-component is zero modulo $E + (\tilde{E} \cap F)$ anyway, and similarly for losing f. Thus the only thing one must check is that the maps are well-defined, since once it is done, we see equally well that these isomorphisms are intertwining operators.

To check that the maps are well-defined, it is enough to deal with the first one, because of the symmetry. First of all, the map

$$\Phi : E + (\tilde{E} \cap \tilde{F}) \longrightarrow (F + (\tilde{E} \cap \tilde{F}))/(F + (E \cap \tilde{F}))$$

mapping $e + \tilde{g}$ to $\tilde{g}$ is well-defined, because if

$$e_1 + \tilde{g}_1 = e_2 + \tilde{g}_2,$$

(with obvious notation) we get

$$\tilde{g}_1 - \tilde{g}_2 = e_2 - e_1 \in E \cap (\tilde{E} \cap \tilde{F}) = E \cap \tilde{F}$$

and hence $\tilde{g}_1 - \tilde{g}_2$ maps to zero in the right-hand quotient modulo $F + (E \cap \tilde{F})$. It is then enough to check that $E + (\tilde{E} \cap F) \subset \ker(\Phi)$ to see that Φ induces the linear map we want, defined on the left-side quotient. But for an element of the type $e + g$ with $g \in \tilde{E} \cap F$, Φ maps it to g modulo $F + (\tilde{F} \cap E)$, which is 0 since $g \in F$!

Example 2.7.4. (1) Let ϱ be semisimple and finite-dimensional, say

$$\varrho \simeq \varrho_1 \oplus \varrho_2 \oplus \cdots \oplus \varrho_k,$$

with ϱ_i irreducible. Then a sequence as above is provided by

$$E_i = \varrho_1 \oplus \cdots \oplus \varrho_i,$$

with $E_i/E_{i-1} \simeq \varrho_i$. It is clear that if we permute the labels i of the ϱ_i, this does not change ϱ, but the sequence changes; however, the quotients are indeed merely permuted.

In that case $n_\pi(\varrho)$ is the number of components ϱ_i which are isomorphic to π.

(2) Consider $k = \mathbf{C}$ and the group

$$G = \left\{ \begin{pmatrix} z & x \\ 0 & 1 \end{pmatrix} \mid z \in \mathbf{C}^\times, \quad x \in \mathbf{C} \right\}$$

with its 2-dimensional representation given by the inclusion in $\mathrm{GL}_2(\mathbf{C}) = \mathrm{GL}(\mathbf{C}^2)$. This representation is not semisimple. With (e_1, e_2) the canonical basis of $\mathbf{C}^2$, one can take $E_1 = \mathbf{C}e_1$, $E_2 = \mathbf{C}^2$; indeed, E_1 is a subrepresentation because

$$\begin{pmatrix} z & x \\ 0 & 1 \end{pmatrix} e_1 = ze_1,$$

and $E_1/E_0 = E_1$ and E_2/E_1 are one-dimensional, hence irreducible. In abstract terms, E_1 is the representation

$$\begin{pmatrix} z & x \\ 0 & 1 \end{pmatrix} \mapsto z \in \mathbf{C}^\times$$

while E_2/E_1 is in fact the trivial representation of G.

Determining the Jordan-Hölder-Noether irreducible components of a representation ϱ can be delicate. At least the following holds:

Lemma 2.7.5. *Let k be a field, G a group and ϱ a finite-dimensional k-representation of G. If ϱ_1 is a finite-dimensional irreducible representation of G and*

$$\mathrm{Hom}_G(\varrho, \varrho_1) \neq 0, \qquad \text{or} \qquad \mathrm{Hom}_G(\varrho_1, \varrho) \neq 0,$$

then ϱ_1 is among the Jordan-Hölder-Noether composition factors of ϱ.

Proof. Both are similar and very intuitive, so we consider here only the case of a non-zero intertwiner $\varrho \xrightarrow{\Phi} \varrho_1$. Let E be the space on which ϱ acts. Because the image is then a non-zero subrepresentation of ϱ_1, which is irreducible, we see that Φ is surjective. Thus $E^1 = \ker(\Phi) \subset E$ is a proper subrepresentation with

$$E/E^1 \simeq \varrho_1.$$

Constructing then successive maximal proper subrepresentations of E^1, say E^2, E^3, ..., we obtain a composition series of ϱ where ϱ_1 is one of the composition factors. By uniqueness, this means ϱ_1 is indeed one of the composition factors of ϱ. □

Example 2.7.6. If a representation $\varrho : G \longrightarrow \mathrm{GL}(E)$ has, for a given basis, a matrix representation which is block triangular

$$\varrho^{\boldsymbol{m}}(g) = \left(\begin{array}{c|c|c|c} \varrho_1^{\boldsymbol{m}}(g) & \star & \dots & \star \\ \hline 0 & \varrho_2^{\boldsymbol{m}}(g) & \star & \dots \\ \hline \vdots & \vdots & & \vdots \\ \hline 0 & 0 & \dots & \varrho_n^{\boldsymbol{m}}(g) \end{array}\right).$$

with square blocks of size $d_1, \ldots, d_n$ on the diagonal, then the $\varrho_i^{\boldsymbol{m}}$ are matrix representations of the composition factors of ϱ. Indeed, multiplication shows that $g \mapsto \varrho_i^{\boldsymbol{m}}(g)$ is a homomorphism to $\mathrm{GL}_{d_i}(k)$, and the subspaces E_i can be defined as those spanned by the

$$d_1 + \cdots + d_i$$

first basis vectors; as in (2.30), one sees that on E_i/E_{i-1}, ϱ acts in the basis formed of the vectors in the i-th block of the given one, like the matrix representation $\varrho_i^{\boldsymbol{m}}$.

If ϱ is semisimple, we can find a decomposition as above with block *diagonal* matrices, and indeed, a block-diagonal decomposition (with irreducible blocks) corresponds to a semisimple representation. However, if the decomposition turns out to be merely block-triangular (with some non-zero off-diagonal blocks), this does *not* mean that the representation is not semisimple! It might just be that the choice of basis was not the best possible.

Here is an example: consider $G = \mathfrak{S}_3$ and the representation

$$\mathfrak{S}_3 \longrightarrow \mathrm{GL}(\mathbf{C}^3)$$

by permutation of the coordinates. The subspace

$$F = \{(x, y, z) \in \mathbf{C}^3 \mid x + y + z = 0\}$$

is a subrepresentation. In terms of the basis $(1, -1, 0)$, $(1, 0, -1)$, $(1, 0, 0)$ of $\mathbf{C}^3$, we see for instance that the action of the cycle (123) is

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & -1 & -1 \\ 0 & 0 & 1 \end{pmatrix}$$

where the third column shows that it is not block-diagonal. This column is given by

$$\varrho((123))(1, 0, 0) = (0, 0, 1) = (1, 0, 0) - (1, 0, -1).$$

However, the representation is semisimple here (taking the third basis vector $(1, 1, 1)$ will lead to a block-diagonal decomposition).

If

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

is a semisimple representation, it is natural to ask to what extent its irreducible subspaces are unique, and more generally, what its subrepresentations look like. Here again there are possible traps. If

$$E = E_1 \oplus \cdots \oplus E_n \tag{2.34}$$

is *a* decomposition into irreducible subspaces, then we know that the isomorphism classes of the E_i, and their multiplicities, are determined by ϱ, up to isomorphism. The actual subspaces E_i, in general, are not so determined; a related fact is that there are usually many more subrepresentations of E than the obvious ones of the form

$$F = \bigoplus_{i \in S} E_i$$

for some subset $S \subset \{1, \ldots, n\}$. Indeed, this is already clear in the case of the trivial group G, where any representation is semisimple, and a decomposition (2.34) is obtained by any choice of a basis of E: if $\dim(E) \geqslant 2$, the vector space E contains many more subspaces than those spanned by finitely many "axes" spanned by basis vectors.

More generally, consider $E = E_1 \oplus E_1$, a direct sum of two copies of E_1. Then we can also write

$$E = F_1 \oplus F_2,$$

with

$$F_1 = \{(v_1, v_1) \mid v_1 \in E_1\}, \quad F_2 = \{(v_1, -v_2) \mid v_1 \in E_1\}$$

(at least if $2 \neq 0$ in k...), and the maps

$$\begin{cases} E_1 \longrightarrow F_1 \\ v \mapsto (v, v) \end{cases}, \qquad \begin{cases} E_1 \longrightarrow F_2 \\ v \mapsto (v, -v) \end{cases}$$

are isomorphisms of representations.

A weaker uniqueness is still valid: in any decomposition (2.34), the direct sum $M(\pi)$ of *all* subspaces E_i on which the action of G is isomorphic to a given irreducible representation π, is independent of the decomposition.

Proposition 2.7.7 (Unicity of isotypic components). *Let G be a group and let k be a field. Let $\varrho : G \longrightarrow \mathrm{GL}(E)$ be a semisimple k-representation of G.*

(1) *Fix an irreducible k-representation π of G. For any decomposition*

$$E = \bigoplus_{i \in I} E_i, \tag{2.35}$$

where ϱ acts irreducibly on the subspaces E_i, the subspace

$$\bigoplus_{E_i \simeq \pi} E_i \subset E$$

is the same. Indeed, it is equal to the sum *of all subrepresentations of E isomorphic to π. This space is called* the π-isotypic component *of E and is denoted $M(\pi)$ or $M_E(\pi)$.*

(2) *In particular, if all irreducible components ϱ_i of ϱ occur with multiplicity 1, the corresponding subspaces $E_i \subset E$ isomorphic to ϱ_i are unique, and any subrepresentation of E is equal to*

$$\bigoplus_{i \in S} E_i$$

for some subset $S \subset I$.

(3) *If ϱ_1, ϱ_2 are semisimple k-representations of G, acting on E_1 and E_2 respectively, and if $\Phi : E_1 \to E_2$ is a G-homomorphism, then the restriction of Φ to the isotypic component $M_{E_1}(\pi)$ is a linear map*

$$M_{E_1}(\pi) \longrightarrow M_{E_2}(\pi),$$

i.e., the image of $M_{E_1}(\pi)$ is contained in $M_{E_2}(\pi)$.

Proof. In order to prove (1), we denote by $M(\pi)$ the sum (not necessarily direct, of course) of the subrepresentations of E isomorphic to π; this is a well-defined intrinsic subspace of E, and it is clear that, for any decomposition (2.35), we have

$$\bigoplus_{E_i \simeq \pi} E_i \subset M(\pi). \tag{2.36}$$

We thus need only prove the converse inclusion. But if $F \subset E$ is a subrepresentation isomorphic to π, and j is such that the representation on E_j is *not* isomorphic to π, the projection map

$$p_j \,:\, F \longrightarrow E_j$$

(defined using (2.35)) is in $\mathrm{Hom}_G(F, E_j) = 0$, by Schur's Lemma 2.2.4. Thus the component along E_j of any vector in F is zero, and that means precisely that F is a subspace of the left-hand side of (2.36). From the definition of $M(\pi)$, and the fact that F was arbitrary, we get the first part of the proposition.

Now the first part of (2) follows from (1), since in the absence of multiplicity $\geqslant 2$, the intrinsic isotypic components are reduced to a single E_i in the decomposition (2.35). And if $F \subset E$ is a subrepresentation, we know (Lemma 2.2.7) that F is also semisimple, and then any of its own irreducible subspace is one in E, and hence is equal to some E_i. Thus F becomes equal to the direct sum of those subspaces E_i which are in F.

Finally, (3) is due to the fact that there exists (from (1)) an intrinsic definition of $M_{E_1}(\pi)$, which must naturally be "transported" under an intertwining map to E_2. Precisely, $M_{E_1}(\pi)$ is generated by the vectors v in the image of all homomorphisms $\Psi \,:\, \pi \longrightarrow \varrho_1$. For any such map, the composite

$$\Phi \circ \Psi \,:\, \pi \longrightarrow \varrho_2$$

has image in $M_{E_2}(\pi)$, for the same reason. This image contains of course the image under Φ of all vectors $v \in \mathrm{Im}(\Psi)$. This means that $M_{E_2}(\pi)$ contains the image under Φ of generators of $M_{E_1}(\pi)$, and this means that the π-isotypic component of E_2 contains the image of $M_{E_1}(\pi)$, which is what the statement of (3). □

Example 2.7.8 (Isotypic components for the trivial and one-dimensional representations). The simplest example concerns the trivial representation $\mathbf{1}_G$. This is always irreducible, and for any representation $\varrho \,:\, G \longrightarrow \mathrm{GL}(E)$, we have

$$M_\varrho(\mathbf{1}) = \varrho^G,$$

the subspace of invariant vectors (this is something that was stated, without using the same words, in Example 2.1.7).

For instance, if ϱ is the regular representation of G, we find (from the same example) that $M_{\mathrm{reg}}(\mathbf{1})$ is the one-dimensional subspace of *constant* k-valued functions on G.

More generally, if $\dim(\pi) = 1$, so that it is automatically irreducible, and $\pi(g) \in k^\times$ is just a scalar, we have

$$M_\varrho(\pi) = \{v \in E \mid \varrho(g)v = \pi(g)v \text{ for all } g \in G\}.$$

Applied to the regular representation, the reader will easily check that $M_{\mathrm{reg}}(\pi)$ is again one-dimensional, and is generated by π itself, seen as a k-valued function.

Example 2.7.9. Here is an application of these ideas, leading to the proof of Theorem 2.6.1. We consider the representation ϱ_m of $G = \mathrm{SL}_2(\mathbf{C})$ on the space V_m of homogeneous polynomials of degree m in $\mathbf{C}[X, Y]$ (Example 2.6.1).

Proposition 2.7.10. *For each $m \geqslant 0$, the representation ϱ_m of $\mathrm{SL}_2(\mathbf{C})$ is irreducible.*

Proof. We use a fairly common strategy, which we will see again later, to attempt to analyze ϱ_m (which, at this point, we do not know to be semisimple): we first consider the restriction to some subgroup of G for which we understand well (or better, at least) the representation theory. Here we consider

$$T = \left\{ t(\lambda) = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix} \mid \lambda \in \mathbf{C}^\times \right\} \simeq \mathbf{C}^\times$$

(a choice justified partly by the fact that T is abelian). We can see easily how T acts on the basis vectors $e_i = X^i Y^{m-i}$, $0 \leqslant i \leqslant m$: for $\lambda \in \mathbf{C}^\times$, we have by definition

$$\varrho_m(t(\lambda))e_i = (\lambda X)^i (\lambda^{-1} Y)^{m-i} = \lambda^{2i-m} e_i.$$

This means that the lines $\mathbf{C}e_i$ are all stable under the action of T, and that $\operatorname{Res}_T^G \varrho_m$ acts on $\mathbf{C}e_i$ according to the representation

$$\chi_{2i-m} \begin{cases} T & \longrightarrow & \mathrm{GL}_1(\mathbf{C}) \\ t(\lambda) & \mapsto & \lambda^{2i-m}. \end{cases}$$

Since (e_i) is a basis of V_m, this means that we have proved that

$$\operatorname{Res}_T^G(\varrho_m) \simeq \bigoplus_{0 \leqslant i \leqslant m} \chi_{2i-m} = \chi_{-m} \oplus \chi_{-m+2} \oplus \cdots \oplus \chi_m. \tag{2.37}$$

The right-hand side is therefore semisimple, and its irreducible components (the χ_{2i-m}, which are irreducible since one-dimensional) occur with multiplicity 1.

Now consider any non-zero G-stable subspace $F \subset V_m$; it is also a subrepresentation of the restriction of ϱ_m to T, obviously, and from what we just observed, Proposition 2.7.7, (2), implies that the subspace F is a direct sum of some of the lines $\mathbf{C}e_i$ corresponding to the representations of T occurring in F. Thus there exists some non-empty subset

$$I \subset \{0, \ldots, m\}$$

such that

$$F = \bigoplus_{i \in I} \mathbf{C}e_i. \tag{2.38}$$

Now we "bootstrap" this information using the action of other elements of G than those in T. Namely, fix some $i \in I$ and consider the action of

$$u = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \; ; \tag{2.39}$$

since F is a stable under G, we know that $\varrho_m(u)e_i \in F$, and this means

$$X^i (X+Y)^{m-i} \in F.$$

Expanding by the binomial theorem, we get

$$X^m + (m-i)X^{m-1}Y + \cdots + (m-i)X^i Y^{m-i-1} + X^i Y^{m-i} = \sum_{j=i}^{m} \binom{m-i}{j-i} e_j \in F,$$

and comparison with (2.38) leads to the conclusion that all $j \geqslant i$ are also in I. Similarly, considering the action of

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

we conclude that $j \in I$ if $j \leqslant i$. Hence, we must have $I = \{0, \ldots, m\}$, which means that $F = V_m$. This gives the irreducibility. □

Exercise 2.7.11 (Irreducibility of ϱ_m restricted to a smaller group). Consider again the representation ϱ_m of $\mathrm{SL}_2(\mathbf{C})$ for $m \geqslant 0$. We restrict it now to the subgroup $\mathrm{SU}_2(\mathbf{C})$ of unitary matrices of size 2. The proof of irreducibility of ϱ_m in the previous example used the element (2.39) and its transpose, which do not belong to $\mathrm{SU}_2(\mathbf{C})$. However, ϱ_m restricted to $\mathrm{SU}_2(\mathbf{C})$ is still irreducible, as claimed in Theorem 2.6.1. Of course, this gives another proof of the irreducibility of ϱ_m as a representation of $\mathrm{SL}_2(\mathbf{C})$.

(1) Show that a decomposition (2.38) still holds with I not empty for a non-zero subspace F stable under $\mathrm{SU}_2(\mathbf{C})$.

(2) Let j be such that $e_j = X^j Y^{n-j}$ is in F. Show that for

$$g = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \in \mathrm{SU}_2(\mathbf{C}),$$

we have

$$\varrho_m(g)e_j = \sum_{0 \leqslant i \leqslant m} f_i(\theta) e_i$$

where the f_i are functions on $[0, 2\pi]$ which are not identically zero. Deduce that $F = V_m$.

Exercise 2.7.12 (Another example of irreducible representation). The following example will be used in Chapter 6. We consider $k = \mathbf{C}$ and we let $V = \mathbf{C}^n$ for $n \geqslant 1$ and $E = \mathrm{End}(V)$. The group $G = \mathrm{GL}_n(\mathbf{C})$ acts on V (by matrix multiplication!) and therefore there is an associated representation on E, as in (2.15).

(1) Show that this representation on E is the conjugation action

$$g \cdot A = gAg^{-1}$$

and that the space E_0 of endomorphisms $A \in E$ with trace 0 is a subrepresentation of E.

We will now prove that E_0 is an irreducible representation of G, and in fact that it is already an irreducible representation of the subgroup $\mathrm{SU}_n(\mathbf{C})$ of unitary matrices with determinant 1.

(2) Let $T \subset \mathrm{SU}_n(\mathbf{C})$ be the diagonal subgroup of $\mathrm{SU}_n(\mathbf{C})$. Show that the restriction of E_0 to T decomposes as the direct sum of T-subrepresentations

$$E_0 = H \oplus \bigoplus_{i \neq j} \mathbf{C} E_{i,j}$$

where H is the subspace of diagonal matrices in E_0 and $E_{i,j} \in E_0$ is, for $1 \leqslant i \neq j \leqslant n$, the rank 1 matrix with a single coefficient equal to 1 on the (i,j)-th entry. Moreover show that the subspaces $\mathbf{C}E_{i,j}$ each carry distinct non-trivial irreducible representations of T, and that $H = (E_0)^T$ is the space of T-invariants.

(3) Let $F \subset E_0$ be a non-zero subspace stable under $\mathrm{SU}_n(\mathbf{C})$. Show that F can not be contained in H.

(4) Deduce that F contains all $E_{i,j}$ for $i \neq j$. Then conclude that $F = E_0$. [Hint: Show that suitable combinations of vectors generating H are $\mathrm{SU}_n(\mathbf{C})$-conjugate of combinations of some $E_{i,j}$, $i \neq j$.]

2.7.2. Schur's Lemma.

The next result is fundamental. It is usually called "Schur's Lemma", but it was known and used by others like Frobenius or Burnside, independently of Schur. In fact, it is a refinement of Lemma 2.2.4; the version we state is valid for finite-dimensional representations, but there are variants when infinite-dimensional representations are considered with some topological restrictions (as we will see later).

Proposition 2.7.13 (Schur's Lemma, II). *Let G be a group and let k be an algebraically closed field, for instance $k = \mathbf{C}$.* (1) *If π_1 and π_2 are irreducible k-representations of G which are non-isomorphic, we have*

$$\mathrm{Hom}_G(\pi_1, \pi_2) = 0.$$

(2) *If π is an irreducible k-representation of G* of finite dimension*, we have*

$$\mathrm{Hom}_G(\pi, \pi) = \mathrm{Hom}_k(\pi, \pi)^G = k\mathrm{Id}_\pi.$$

(3) *Conversely, if π is a finite-dimensional,* semisimple *k-representation of G such that* $\dim \mathrm{Hom}_G(\pi, \pi) = 1$*, it follows that π is irreducible.*

Note that we used here the natural representation of G on homomorphism spaces, in which the G-homomorphisms are the G-invariants. The statement gives a very strong expression of the fact that non-isomorphic irreducible representations of G are "independent" of each other; it is frequently used in the form of the formula

$$\dim \mathrm{Hom}_G(\pi_1, \pi_2) = \delta(\pi_1, \pi_2) = \begin{cases} 1 & \text{if } \pi_1 \simeq \pi_2 \\ 0 & \text{otherwise,} \end{cases} \tag{2.40}$$

for irreducible finite-dimensional representations of a group G over an algebraically closed field.

We will see other incarnations of this fact later (e.g., Theorem 2.7.24).

Proof. The first part is a consequence of the earlier version of Schur's Lemma (Lemma 2.2.4). For the second, consider a G-homomorphism Φ from π to itself. The fact that k is algebraically closed and π is finite-dimensional implies that Φ, as a linear map, has an eigenvalue, say $\lambda \in k$. But $\Phi - \lambda\mathrm{Id}$ is then a G-homomorphism of π which is *not injective.* By (2), the only possibility is that $\Phi - \lambda\mathrm{Id}$ be identically zero, which is the desired conclusion.

Finally we prove the converse when π is semisimple. Let E be the k-vector space on which π acts; we can assume that $\dim E \geqslant 1$, and then we let $F \subset E$ be an irreducible subrepresentation, and F_1 a complementary subrepresentation, so that $E = F \oplus F_1$. The projector $\Phi : E \longrightarrow E$ onto F with kernel F_1 is an element of $\mathrm{Hom}_G(\pi, \pi)$ (we saw this explicitly in Lemma 2.2.2), and our assumption on π implies that it is a multiple of the identity. Since it is non-zero, it is therefore equal to the identity, which means that $F_1 = 0$ and $E = F$ is irreducible. □

Exercise 2.7.14 (Schur's Lemma and semisimplicity). The last statement in Schur's Lemma can be a very useful irreducibility criterion. However, one should not forget the semisimplicity condition! Consider the representation ϱ of

$$G = \left\{ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \right\} \subset \mathrm{GL}_2(\mathbf{C})$$

on $\mathbf{C}^2$ by left-multiplication.

(1) What are its composition factors? Is it semisimple?

(2) Compute $\mathrm{Hom}_G(\varrho, \varrho)$ and conclude that the converse of Schur's Lemma (part (3)) does not always hold when π is not assumed to be semisimple. What happens if instead of G one uses its subgroup where $a = d = 1$?

A simple, but important, corollary of Schur's Lemma is the following:

Corollary 2.7.15 (Abelian groups and central character). *Let G be an abelian group, k an algebraically closed field. Then any finite-dimensional irreducible representation of G is of dimension* 1*, i.e., the finite-dimensional irreducible representations of G coincide with the homomorphisms $G \longrightarrow k^{\times}$.*

More generally, if G is any group, and $\varrho : G \longrightarrow \mathrm{GL}(E)$ is a finite-dimensional irreducible representation of G, there exists a one-dimensional representation ω of the center $Z(G)$ of G such that

$$\varrho(z) = \chi(z)\mathrm{Id}_E$$

for all $z \in Z(G)$. This representation ω is called the central character *of ϱ.*

Proof. (1) Let ϱ be a finite-dimensional irreducible representation of G, acting on E. *Because G is abelian*, any $\Phi = \varrho(g) : E \to E$ is in fact a homomorphism in $\mathrm{Hom}_G(\varrho, \varrho)$. By Schur's Lemma 2.7.13, there exists therefore $\lambda(g) \in k$ such that $\varrho(g) = \lambda(g)\mathrm{Id}$ is a scalar. Then any one-dimensional subspace of E is invariant under all operators $\varrho(g)$, and by irreducibility, this means that E is equal to any such subspace.

(2) Similarly, for G arbitrary, if z is an element of the center of G, we see that $\varrho(z)$ commutes with all $\varrho(g)$, for any representation of G, i.e., $\varrho(z) \in \mathrm{End}_G(\varrho)$. If ϱ is irreducible, Schur's Lemma implies that $\varrho(z)$ is multiplication by a scalar, and of course the latter is a one-dimensional representation of $Z(G)$. □

Remark 2.7.16 (Division algebras). Example 2.4.1 shows that this result does not hold in general if the field is not necessarily algebraically closed.

If ϱ is an irreducible (finite-dimensional) k-representation of G, the earlier version Schur's Lemma already shows that $A = \mathrm{End}_G(\varrho)$, the space of G-endomorphisms of ϱ, has a remarkable structure: it is a subalgebra of the matrix algebra $\mathrm{End}_k(\varrho)$ which is a *division algebra*, i.e., any non-zero element of A has an inverse in A.

In the case of Example 2.4.1, the reader is invited to show explicitly that A is isomorphic to $\mathbf{C}$, as an $\mathbf{R}$-algebra.

Another easy useful corollary is the following algebraic characterization of multiplicities of irreducible representations in a semisimple representation:

Corollary 2.7.17 (Multiplicities). *Let G be a group and k an algebraically closed field. If ϱ is a finite-dimensional semisimple k-representation of G, then for any irreducible k-representation π of G, we have*

$$n_{\pi}(\varrho) = \dim \mathrm{Hom}_G(\varrho, \pi) = \dim \mathrm{Hom}_G(\pi, \varrho),$$

where $n_{\pi}(\varrho)$ is the multiplicity of π as a summand in ϱ.

Proof. If we express

$$\varrho \simeq \bigoplus_i \varrho_i,$$

where ϱ_i are (necessarily finite-dimensional) irreducible representations of G, then we have

$$\mathrm{Hom}_G(\pi, \varrho) \simeq \bigoplus_i \mathrm{Hom}_G(\pi, \varrho_i)$$

for all irreducible representation π. This space has dimension equal to the number of indices i for which $\varrho_i \simeq \pi$, by Schur's Lemma (i.e., by (2.40)), which is of course $n_{\pi}(\varrho)$. A similar argument applies of course to $\mathrm{Hom}_G(\varrho, \pi)$. □

If k is algebraically closed, we can also use Schur's Lemma to give a nice description of the isotypic component $M(\pi)$ of a finite-dimensional semisimple representation ϱ of G (acting on E). To describe this, let E_π be the space on which π acts; then there is a natural k-linear map

$$\Theta \begin{cases} \mathrm{Hom}_G(E_\pi, E) \otimes E_\pi & \longrightarrow & E \\ \Phi \otimes v & \mapsto & \Phi(v). \end{cases}$$

The image of this map is, almost by definition, equal to the isotypic component $M(\pi) \subset E$ (because any non-zero $\Phi \in \mathrm{Hom}_G(E_\pi, E)$ is injective by Schur's Lemma, so that $\Phi(v)$ is in the subrepresentation $\mathrm{Im}(\Phi)$ isomorphic to ϱ.)

If E is finite-dimensional, we then see (by the previous corollary) that the dimensions of $M(\pi)$ and of the source

$$\mathrm{Hom}_G(E_\pi, E) \otimes E_\pi$$

coincide, and we conclude that Θ gives an isomorphism

$$\mathrm{Hom}_G(E_\pi, E) \otimes E_\pi \simeq M(\pi) \subset E.$$

Moreover, Θ is a G-homomorphism if we let G act trivially on $\mathrm{Hom}_G(E_\pi, E)$ (which is natural by (2.18)) and through π on E_π. From this (picking, if needed, a basis of $\mathrm{Hom}_G(E_\pi, E)$) we see that $M(\pi)$ is isomorphic, as representation of G, to a direct sum of d copies of π, where

$$d = \dim \mathrm{Hom}_G(E_\pi, E).$$

As it happens, the injectivity of Θ can also be proved directly, and this leads to the following useful result, where ϱ is not assumed to be semisimple, but we can still characterize the π-isotypic component using the same construction:

Lemma 2.7.18 (A formula for isotypic components). *Let G be a group, and let k be an algebraically closed field. If*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

is a finite-dimensional k-representation of G and

$$\pi : G \longrightarrow \mathrm{GL}(E_\pi)$$

is any irreducible k-representation, the map

$$\Theta \begin{cases} \mathrm{Hom}_G(E_\pi, E) \otimes E_\pi & \longrightarrow & E \\ \Phi \otimes v & \mapsto & \Phi(v) \end{cases}$$

is injective and its image is equal to the π-isotypic component $M_E(\pi)$, the sum of all subrepresentations of E isomorphic to π. In particular,

$$M(\pi) \simeq (\dim \mathrm{Hom}_G(E_\pi, E))\pi$$

as representation.

Note that, if ϱ is not semisimple, π might also appear as composition factor outside of $M(\pi)$ (i.e., as a genuine quotient or subquotient.)

Proof. As before, it is clear that the image of Θ is equal to the subrepresentation $M(\pi) \subset E$. It remains thus to show that Θ is injective, as this leads to the isomorphism

$$M(\pi) \simeq \mathrm{Hom}_G(E_\pi, E) \otimes \pi,$$

from which the last step follows.

Let (Φ_j) be a basis of the space $\mathrm{Hom}_G(E_\pi, E)$, so that any element of the tensor product is of the form

$$\sum_j \Phi_j \otimes v_j$$

for some $v_j \in E_\pi$. Then we have

$$\Theta\Bigl(\sum_j \Phi_j \otimes v_j\Bigr) = \sum_j \Phi_j(v_j) \in E,$$

and the injectivity of Θ is seen to be equivalent to saying that the spaces $F_j = \mathrm{Im}(\Phi_j)$ are in direct sum in E. We prove this in a standard manner as follows: assume the contrary is true, and let $J \subset I$ be a set of smallest order for which there is a relation

$$\sum_{j \in J} \Phi_j(v_j) = 0$$

with $\Phi_j(v_j) \neq 0$ for $j \in J$. Consider any $\ell \in J$; we find that

$$\Phi_\ell(v_\ell) \in \mathrm{Im}(\Phi_\ell) \cap \bigoplus_{j \neq \ell} \mathrm{Im}(\Phi_j) \neq 0,$$

so that, by irreducibility, this intersection is in fact equal to $\mathrm{Im}(\Phi_\ell)$ (note that we wrote that the Φ_j, $j \neq \ell$, are in direct sum because otherwise we could replace the set J by $J - \{\ell\}$.) This means that Φ_ℓ belongs, in an obvious sense, to the space

$$\mathrm{Hom}_G(E_\pi, \bigoplus_{j \neq \ell} \mathrm{Im}(\Phi_j))$$

which is spanned by the homomorphisms $(\Phi_j)_{j \neq \ell}$. This is impossible however, since all the Φ_j's are linearly independent. □

The following addition is also useful:

Lemma 2.7.19. *With assumptions and notation as in Lemma 2.7.18, if (π_i) is a family of pairwise non-isomorphic irreducible representations of G, the isotypic subspaces $M(\pi_i) \subset E$ are in direct sum.*

Proof. Indeed, for any fixed π in this family, the intersection of $M(\pi)$ with the sum of all $M(\pi_i)$, $\pi_1 \neq \pi$, is necessarily zero: it can not contain any irreducible subrepresentation, since the possibilities coming from π are incompatible with those coming from the other π_i. □

2.7.3. Burnsides's irreducibility criterion and its generalizations, 1. We now show how to use Schur's Lemma to prove a result of Burnside which provides a frequently useful irreducibility criterion for finite-dimensional representations, and we derive further consequences along the same lines. In fact, we will prove this twice (and a third time in Chapter 4 in the case of finite groups); in this section, we argue in the style of Burnside, and below in Section 2.7.4, we will recover the same results in the style of Frobenius.

We will motivate the first result, Burnside's criterion, from the following point of view: given a finite-dimensional k-representation ϱ of a group G, acting on the vector space E, the image of ϱ is a subset of the vector space $\mathrm{End}_k(E)$. We ask then "what are the linear relations satisfied by these elements?" For instance, the block-diagonal shape (2.30) of a representation which is is not irreducible shows clearly *some* relations: those that express that the matrices in the bases indicated have lower-left corner(s) equal to 0, for instance. These are obvious. Are there others?

Theorem 2.7.20 (Burnside's irreducibility criterion). *Let k be an algebraically closed field, G a group. A* finite-dimensional *k-representation*

$$\varrho\,:\,G\longrightarrow \mathrm{GL}(E)$$

is irreducible *if and only if the image of ϱ satisfies no non-trivial linear relation in* $\mathrm{End}_k(E)$*. Equivalently, ϱ is irreducible if and only if the linear span of the image of ϱ in* $\mathrm{End}_k(E)$ *is* equal *to* $\mathrm{End}_k(E)$*.*

The proof we give is a modern version of Burnside's original argument. One can give much shorter proofs – the one in the next section is an example, – but this one has the advantage of "exercising" the basic formalism of representation theory, and of being easy to motivate.

Proof. First of all, the two statements we give are equivalent by duality of finite-dimensional k-vector spaces. More precisely, let $V=\mathrm{End}_k(E)$; then by "relations satisfied by the image of ϱ", we mean the k-linear subspace

$$R=\{\phi\in V'\mid\langle\phi,\varrho(g)\rangle=0\text{ for all }g\in G\},$$

of V', the linear dual of $\mathrm{End}_k(E)$. Then we are saying that $R=0$ if and only if the image of G spans V, which is part of duality theory.

The strategy of the proof is going to be the following:

(1) For some natural representation of G on V', we show that R is a subrepresentation;
(2) We find an explicit decomposition of V' (with its G-action) as a direct sum of irreducible representations, embedded in V' in a specific manner;
(3) Using this description, we can see what the possibilities for R are, and in particular that $R=0$ if ϱ is irreducible.

This strategy will also be used afterward to give a more general result of comparison of distinct irreducible representations.

We let G act on V' by the contragredient of the representation of G on V given by

$$g\cdot T=\varrho(g)\circ T$$

for $g\in G$ and $T\,:\,E\to E$. Note that this corresponds to the action (2.19), and not the action (2.16). To check that $R\subset V'$ is a subrepresentation, we need simply note that if $\phi\in R$ and $g\in G$, then we have

$$\langle g\cdot\phi,\varrho(h)\rangle=\langle\phi,g^{-1}\cdot\varrho(h)\rangle=\langle\phi,\varrho(g^{-1}h)\rangle=0$$

for all $h\in G$, which means that $g\cdot\phi$ is also in R.

From (2.20), we know that V – with the above action – is isomorphic to a direct sum of $\dim\varrho$ copies of ϱ; hence V' is isomorphic to a direct sum of the same number of copies of $\tilde{\varrho}$, which we know to be irreducible (Lemma 2.2.13). It follows that any irreducible subrepresentation π of R – which exists if $R\neq 0$ – must be itself isomorphic to $\tilde{\varrho}$. (This fact is clear from the Jordan-Hölder-Noether Theorem, but can be seen directly also by considering the composites

$$p_i\,:\,\pi\hookrightarrow R\hookrightarrow V'\simeq\bigoplus_{i\leqslant\dim\varrho}W_i\longrightarrow W_i\simeq\tilde{\varrho},$$

where W_i are subrepresentations of V' isomorphic to $\tilde{\varrho}$; each of these composites is in $\mathrm{Hom}_G(\pi,\tilde{\varrho})$, and hence is either 0 or an isomorphism, by Schur's Lemma. Not all p_i can be zero, if $\pi\neq 0$, hence $\pi\simeq\tilde{\varrho}$.)

However, we claim that there is an isomorphism (of k-vector spaces)

$$\begin{cases} E & \longrightarrow & \mathrm{Hom}_G(E', V') \\ v & \mapsto & \alpha_v \end{cases}$$

where $\alpha_v : E' \to V'$ is defined by

$$\langle \alpha_v(\lambda), T \rangle = \langle \lambda, T(v) \rangle$$

for $\lambda \in E'$ and $T : E \to E$. If this is the case, then assuming that $R \neq 0$, and hence that R contains a copy of $\tilde{\varrho}$, means that, for some $v \neq 0$, the image of α_v is in R. But this implies that for all $\lambda \in E'$, and $g \in G$, we have

$$0 = \langle \alpha_v(\lambda), \varrho(g) \rangle = \langle \lambda, \varrho(g)v \rangle$$

which is impossible even for a single g, since $\varrho(g)v \neq 0$.

Checking the claim is not very difficult; we leave it to the reader to verify that each α_v is indeed a G-homomorphism from E' to V', and that $v \mapsto \alpha_v$ is k-linear. We then observe that

$$\dim \mathrm{Hom}_G(E', V') = \dim \mathrm{Hom}_G(\tilde{\varrho}, (\dim E)\tilde{\varrho}) = \dim(E) \dim_G(\tilde{\varrho}, \tilde{\varrho}) = \dim E$$

by Schur's Lemma again (using the fact that k is algebraically closed). So the map will be an isomorphism as soon as it is injective. However, $\alpha_v = 0$ means that

$$\langle \lambda, T(v) \rangle = 0$$

for all $\lambda \in E'$ and $T \in V$, and that only happens when $v = 0$ (take T to be the identity). □

Example 2.7.21. Consider again the representation ϱ of Example 2.4.1 which is irreducible over $\mathbf{R}$ but not absolutely irreducible. We see then that the linear span of the image of ϱ is the proper subalgebra

$$\left\{ \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \mid a,\ b \in \mathbf{R} \right\}$$

in $\mathrm{M}_2(\mathbf{R})$. In particular, it does satisfy non-trivial relations like "the diagonal coefficients are equal".

We emphasize again the strategy we used, because it is in fact a common pattern in applications of representation theory: one wishes to analyze a certain vector space (here, the relation space R); this space is seen to be a subspace of a bigger one, on which a group G acts, and then the space is seen to be a *subrepresentation* of this bigger space; independent analysis of the latter is performed, e.g., a decomposition in irreducible summands; and then one deduces *a priori* restrictions on the possibilities for the space of original interest.

We now implement this again in a generalization of Burnside's Theorem (which is due to Frobenius and Schur). To motivate it, consider a finite-dimensional k-representation of G

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

which is irreducible, with k algebraically closed. Burnside's Theorem means, in particular, that if we fix a basis (v_i) of E and express ϱ in the basis (v_i), producing the homomorphism

$$\varrho^{\boldsymbol{m}} : G \longrightarrow \mathrm{GL}(\dim(E), k)$$

the resulting "matrix coefficients" $(\varrho^{\boldsymbol{m}}_{i,j}(g))$, seen as functions on G, are k-linearly independent. Indeed, if we denote by (λ_i) the dual basis of E', we have

$$\varrho^{\boldsymbol{m}}_{i,j}(g) = \langle \lambda_i, \varrho(g)v_j \rangle,$$

so that a relation

$$\sum_{i,j} \alpha_{i,j}\varrho^{\boldsymbol{m}}_{i,j}(g) = 0$$

valid for all g, for some fixed $\alpha_{i,j} \in k$, means that the element ϕ of $\mathrm{End}_k(E)'$ defined by

$$\langle \phi, T \rangle = \sum_{i,j} \alpha_{i,j} \langle \lambda_i, T(v_j) \rangle$$

is in the relation space R of the proof above, hence is zero.

The interest of these matrix coefficients is that they are functions on G (with values in k); as such, they might be written down without mentioning at all the representation, and in particular the representation space. However, the choice of basis is annoying, so the following definition is typically better:

Definition 2.7.22 (Matrix coefficient). Let G be a group, k a field and

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

a k-representation of G. A *matrix coefficient* of ϱ is any function on G of the type

$$f_{v,\lambda} \begin{cases} G \to k \\ g \mapsto \lambda(\varrho(g)v) = \langle \lambda, \varrho(g)v \rangle, \end{cases}$$

for some fixed $v \in E$ and $\lambda \in E'$.

Remark 2.7.23. Note that these functions are typically *not* multiplicative. An exception is when $\varrho : G \longrightarrow \mathrm{GL}_1(k) = k^\times$ is a one-dimensional representation; in that case one can take $v = 1 \in k$ and λ the identity of k, so that the matrix coefficient is equal to ϱ as a function $G \longrightarrow k$.

Now we come back to the discussion: matrix coefficients of a fixed irreducible representation are k-linearly independent. What is more natural than to ask, "What about different representations?" Is it possible that their matrix coefficients satisfy non-trivial linear relations? The answer is very satisfactory: No! This is another expression of that the fact that distinct (i.e., non-isomorphic) irreducible representations of a group are "independent".

Theorem 2.7.24 (Linear independence of matrix coefficients). *Let G be a group, k an algebraically closed field.*

(1) *For any finite collection (ϱ_i) of pairwise non-isomorphic, finite-dimensional, irreducible k-representations of G, acting on E_i, let*

$$\varrho = \bigoplus_i \varrho_i \text{ acting on } E = \bigoplus_i E_i.$$

Then the k-linear span of the elements $\varrho(g)$, for $g \in G$, in $\mathrm{End}_k(E)$ is equal to

$$\bigoplus_i \mathrm{End}_k(E_i). \tag{2.41}$$

(2) *The matrix coefficients of finite-dimensional irreducible k-representations of G are linearly independent, in the following sense: for any finite collection (ϱ_i) of pairwise non-isomorphic, finite-dimensional, irreducible k-representations of G, acting on E_i, for any choice $(v_{i,j})_{1\leqslant j\leqslant \dim E_i}$ of bases of E_i, and for the dual bases $(\lambda_{i,j})$, the family of functions*

$$(f_{v_{i,j},\lambda_{i,k}})_{i,j,k}$$

on G are k-linearly independent.

Note that there are

$$\sum_i (\dim E_i)^2$$

functions on G in this family, given by

$$\begin{cases} G \to k \\ g \mapsto \langle \lambda_{i,j}, \varrho_i(g)v_{i,k}\rangle_{E_i}, \end{cases}$$

for $1 \leqslant j,k \leqslant \dim E_i$.

Proof. It is easy to see first that (1) implies (2): writing down as above a matrix representation for $\varrho = \bigoplus \varrho_i$ in the direct sum of the given bases of E_i, a k-linear relation on the matrix coefficients implies one on the k-linear span of the image of ϱ, but there is no non-trivial such relation on

$$\bigoplus_i \mathrm{End}_k(E_i).$$

To prove (1), we use the same strategy as in the proof of Burnside's theorem, for the representation ϱ of G on E: let $V = \mathrm{End}_k(E)$ and

$$R = \{\phi \in V' \mid \langle \phi, \varrho(g)\rangle = 0 \text{ for all } g \in G\}$$

be the relation space. We will compute R and show that $R = S$ where

$$S = \{\phi \in V' \mid \langle \phi, T\rangle = 0 \text{ for all } T \in \bigoplus_i \mathrm{End}_k(E_i)\}, \tag{2.42}$$

so that by duality, the linear span of $\varrho(g)$ is equal to (2.41), as claimed.

As before, we consider the representation of G on V' by the contragredient of the action $g \cdot T = \varrho(g) \circ T$ on V, and we see that $R \subset V'$ is a subrepresentation, and that $S \subset V'$ is also one (because ϱ leaves each E_i stable).

We now show that V' is semisimple and exhibit a decomposition into irreducibles. For this purpose, denoting

$$V_{i,j} = \mathrm{Hom}_k(E_j, E_i),$$

we have also (as in Remark 2.2.15) the similar action of G on each $V_{i,j}$, and we obtain first – as k-vector spaces – the direct sum decompositions

$$V = \bigoplus_{i,j} V_{i,j}, \qquad V' = \bigoplus_{i,j} V'_{i,j}$$

(rigorously, we identify $V_{i,j}$ with the subspace of V consisting of all $T : E \to E$ that map the summand E_j to E_i and all other E_ℓ's to 0; this identification is, in fact, also implicit in the statement of the theorem involving (2.41)). These subspaces $V'_{i,j}$ are not irreducible in general: by (2.20), we get isomorphisms of representations

$$V_{i,j} \simeq (\dim E_j)E_i,$$

and hence

$$V'_{i,j} \simeq (\dim E_j)\tilde{\varrho}_i,$$

for the contragredient, leading to the decomposition

$$V' \simeq \bigoplus_{i,j} (\dim E_j)\tilde{\varrho}_i \simeq \bigoplus_{i} (\dim E)\tilde{\varrho}_i,$$

of V' as direct sum of irreducible representations.

Since R and $S \subset V'$ are subrepresentations, they are therefore also semisimple and have irreducible components among the $\tilde{\varrho}_i$. We now determine which subspaces of V', isomorphic to some $\tilde{\varrho}_i$, can be in R.

Fix an index i. As in the proof of Burnside's Theorem, we first claim there is again an isomorphism of k-vector spaces

$$\begin{cases} E & \longrightarrow & \mathrm{Hom}_G(E_i', V') \\ v & \mapsto & \alpha_v \end{cases}$$

defined by the formula

$$\langle \alpha_v(\lambda), T\rangle = \langle \lambda, T(v)\rangle$$

for $\lambda \in E_i'$ and $T \in V$, where $\lambda \in E_i'$ is extended to E by being 0 on the other summands E_l, $l \neq j$.

Indeed, the α_v are G-morphisms, and this map is injective (the arguments are the same here as in the case of Burnside's Theorem); then we find that

$$\begin{aligned} \dim \mathrm{Hom}_G(E_i', V') &= \sum_j \dim \mathrm{Hom}_G(\tilde{\varrho}_i, (\dim E)\tilde{\varrho}_j) \\ &= \dim(E) \sum_j \dim \mathrm{Hom}_G(\tilde{\varrho}_i, \tilde{\varrho}_j) = \dim(E) \end{aligned}$$

since, using Schur's Lemma,[14] only the term $j = i$ contributes a non-zero factor 1 to the sum.

Any subspace of V' isomorphic to the fixed E_i' is therefore of the form $\mathrm{Im}(\alpha_v)$ for some $v \in E$. Now $\mathrm{Im}(\alpha_v) \subset R$ is equivalent with

$$\langle \alpha_v(\lambda), \varrho(g)\rangle = 0$$

for all $g \in G$ and $\lambda \in E_i'$. But since ϱ is the direct sum of the ϱ_i and $\lambda \in E_i'$, we have

$$\langle \alpha_v(\lambda), \varrho(g)\rangle = \langle \lambda, \varrho(g)v\rangle = \langle \lambda, \varrho_i(g)v_i\rangle$$

where v_i is the component of v in E_i. Hence (putting $g = 1$) the condition $\mathrm{Im}(\alpha_v) \subset R$ is equivalent with $v_i = 0$.

But a similar computation shows that $\mathrm{Im}(\alpha_v) \subset S$ is *also* equivalent with $v_i = 0$ (see (2.42)). Varying i, we see that R and S contain exactly the same irreducible subrepresentations. Hence $R = S$, and we saw at the beginning that this implies the conclusion (2.41). □

Example 2.7.25 (Linear independence of one-dimensional representations). Following on Remark 2.7.23, since any one-dimensional representation is irreducible, and two of them are isomorphic if and only if they coincide as functions $G \longrightarrow k^\times$, the theorem shows that any family of homomorphisms

$$\chi_i : G \longrightarrow k^\times \subset k,$$

is k-linearly independent in $C_k(G)$ when k is algebraically closed. But in fact, this last assumption is not needed, because we can see the representations χ_i as taking values in

[14] This is the crucial point, where the "independence" of distinct irreducible representations comes into play.

an arbitrary algebraic closure of k, and they remain irreducible when seen in this manner, like all 1-dimensional representations.

This result is important in Galois theory. As one might expect, it is possible to prove it directly, and the reader should attempt to do it (see, e.g., [**26**, Th. VI.4.1]).

The linear independence of matrix coefficients turns out to have many important applications. In particular, it gives quite precise information on the structure of the regular representation of G acting on the space $C_k(G)$ of k-valued functions on the group.

Corollary 2.7.26 (Matrix coefficients as subrepresentations of the regular representation). *Let G be a group, k an algebraically closed field, and ϱ a finite-dimensional irreducible k-representation of G. Let $M(\varrho)$ be the subspace of $C_k(G)$* spanned *by all matrix coefficients $f_{v,\lambda}$ of ϱ.*

(1) *The space $M(\varrho)$ depends only on ϱ up to isomorphism.*

(2) *It is a subrepresentation of the regular representation of G acting on $C_k(G)$; moreover $M(\varrho)$ is semisimple and isomorphic to a direct sum of* $\dim(\varrho)$ *copies of ϱ.*

(3) *Any subrepresentation of $C_k(G)$ isomorphic to ϱ is contained in the subspace $M(\varrho)$, i.e., $M(\varrho)$ is the ϱ-isotypic component of $C_k(G)$, as defined in Lemma 2.7.18.*

For one-dimensional representations ϱ (e.g., $\varrho = \mathbf{1}$), we already computed $M(\varrho)$ in Example 2.1.7: it is a one-dimensional space, spanned by ϱ seen as a k-valued function. This verifies (3) directly in these simple cases.

Proof. We first check (1), which states that $M(\varrho)$ is a canonical subspace of $C_k(G)$. Let E be the space on which ϱ acts and let $\tau : G \longrightarrow \mathrm{GL}(F)$ be a k-representation isomorphic to ϱ, with the linear map

$$\Phi : E \longrightarrow F$$

giving this isomorphism. Then for any $w \in F$ and $\lambda \in F'$, writing $w = \Phi(v)$ for some $v \in E$, we have

$$\begin{aligned} f_{w,\lambda}(g) &= \langle \lambda, \tau(g)w\rangle_F \\ &= \langle \lambda, \tau(g)\Phi(v)\rangle_F \\ &= \langle \lambda, \Phi(\varrho(g)v)\rangle_F \\ &= \langle {}^t\Phi(\lambda), \varrho(g)v\rangle_E = f_{v,{}^t\Phi(\lambda)}(g), \end{aligned}$$

for all $g \in G$, showing that any matrix coefficient for τ is also one for ϱ. By symmetry, we see that $M(\varrho)$ and $M(\tau)$ are *equal* subspaces of $C_k(G)$.

We next check that $M(\varrho) \subset C_k(G)$ is indeed a subrepresentation: for $v \in E$ (the space on which G acts), $\lambda \in E'$, and $g \in G$, we have

$$\mathrm{reg}(g)f_{v,\lambda}(x) = f_{v,\lambda}(xg) = \langle \lambda, \varrho(xg)v\rangle = f_{\varrho(g)v,\lambda}(x).$$

But this formula says more: it also shows that, for a *fixed* $\lambda \in E'$, the linear map

$$\Phi_\lambda : v \mapsto f_{v,\lambda}$$

is an intertwining operator between ϱ and $M(\varrho)$.

Fix a basis (v_j) of the space E on which ϱ acts, and let (λ_j) be the dual basis. By construction, $M(\varrho)$ is spanned by the matrix-coefficients

$$f_{i,j} = f_{v_i,\lambda_j}, \qquad 1 \leqslant i, j \leqslant \dim(\varrho)$$

and from the linear independence of matrix coefficients, these functions form in fact a *basis* of the space $M(\varrho)$. In particular, we have $\dim M(\varrho) = \dim(\varrho)^2$. Now the intertwining operator

$$\Phi = \bigoplus_i \Phi_{\lambda_i} \,:\, \bigoplus_i \varrho \longrightarrow M(\varrho)$$

is surjective (since its image contains each basis vector $f_{i,j}$), and both sides have the same dimension. Hence it must be an isomorphism, which shows that $M(\varrho)$ is isomorphic to $(\dim(\varrho))\varrho$ as a representation of G.

There only remains to check the last part. Let $E \subset C_k(G)$ be a subrepresentation of the regular representation which is isomorphic to ϱ. To show that $E \subset M(\varrho)$, we will check that the elements $f \in E$, which are functions on G, are all matrix coefficients of E: let $\delta \in C_k(G)'$ be the linear form defined by

$$\delta(f) = f(1)$$

for $f \in C_k(G)$. Consider the element $\delta_E \in E'$ which is the restriction of δ to E. Then, for any function $f \in E$ and $x \in G$, we have, by definition of the regular representation

$$\langle \delta_E, \mathrm{reg}(x)f\rangle_E = \mathrm{reg}(x)f(1) = f(x).$$

The left-hand side (as a function of x) is a matrix coefficient for ϱ (since reg on E is isomorphic to ϱ), and hence we see that $f \in M(\varrho)$. □

The next corollary will be improved in the chapter on representations of finite groups. We state it here because it is the first a priori restriction we have found on irreducible representations for certain groups:

Corollary 2.7.27. *Let G be a finite group and let k be an algebraically closed field. There are only finitely many irreducible k-representations of G up to isomorphism, and they satisfy*

$$\sum_{\varrho} (\dim \varrho)^2 \leqslant |G|$$

where the sum is over such isomorphism classes of irreducible k-representations of G.

Proof. First of all, since the space of an irreducible representation of a finite group is spanned by finitely many vectors $\varrho(g)v$, $g \in G$ (for any vector $v \neq 0$), any irreducible representation of G is finite-dimensional.

Then by the previous corollary, for any such irreducible representation ϱ, the regular representation reg_G contains a subspace isomorphic to $\dim(\varrho)$ copies of ϱ. By the linear independence of matrix coefficients of non-isomorphic irreducible representations (Theorem 2.7.24), the sum over ϱ of these representations is a direct sum and has dimension

$$\sum_{\varrho} (\dim \varrho)^2,$$

hence the result. □

Remark 2.7.28. If G is finite, there is equality in this formula if and only if the regular representation is semisimple. (If there is equality, this means that

$$C_k(G) = \bigoplus_{\varrho} M(\varrho) \simeq \bigoplus_{\varrho} (\dim \varrho)\varrho$$

is semisimple: conversely, if $C_k(G)$ is semisimple, by Lemma 2.2.8 there exists a stable subspace F such that

$$C_k(G) = F \oplus \Bigl(\bigoplus_{\varrho} M(\varrho)\Bigr),$$

but F can not contain any irreducible subrepresentation, as it would be isomorphic to some ϱ and hence contained in $M(\varrho)$ by what we have seen.) In Chapter 4, we will see that this semisimplicity occurs if and only if the characteristic of the field k does not divide the order of G. Readers may enjoy trying to think about it beforehand, and should also write down explicitly a matrix representation of (say) the regular representation of $G = \mathbf{Z}/2\mathbf{Z}$ over a field of characteristic 2, to check that the latter is not semisimple.

In the next section, we will derive further consequences of the linear independence of matrix coefficients, related to *characters* of finite-dimensional representations. Before, as another application of Schur's Lemma and its corollaries, we can now prove Proposition 2.3.17 about irreducible representations of a direct product $G = G_1 \times G_2$.

Proof of Proposition 2.3.17. Recall that we are considering a field k and two groups G_1 and G_2 and want to prove that all finite-dimensional irreducible k-representations of $G = G_1 \times G_2$ are of the form $\varrho \simeq \varrho_1 \boxtimes \varrho_2$ for some irreducible representations ϱ_i of G_i (unique up to isomorphism). We will prove the result here when k is algebraically closed, and leave the general case to an exercise for the reader below.

To begin with, if ϱ_1 and ϱ_2 are finite-dimensional irreducible representations of G_1 and G_2, the irreducibility of $\varrho = \varrho_1 \boxtimes \varrho_2$ follows from Burnside's irreducibility criterion: since the $\varrho_1(g_1)$ and $\varrho_2(g_2)$, for $g_i \in G_i$, span the k-linear endomorphism spaces of their respective spaces, it follows by elementary linear algebra that the $\varrho_1(g_1) \otimes \varrho_2(g_2)$ also span the endomorphism space of the tensor product. (Here we use the fact that k is algebraically closed.)

Thus what matters is to prove the converse. Let therefore

$$\varrho : G_1 \times G_2 \longrightarrow \mathrm{GL}(E)$$

be an irreducible k-representation. We restrict ϱ to the subgroup $G_1 = G_1 \times \{1\} \subset G$, and we let $E_1 \subset E$ be an irreducible subrepresentation of E seen as representation of G_1; we denote by ϱ_1 the corresponding "abstract" representation of G_1. We now proceed to find a representation ϱ_2 of G_2 such that

$$\varrho_1 \boxtimes \varrho_2 \simeq \varrho.$$

For this purpose, define the k-vector space

$$E_2 = \mathrm{Hom}_{G_1}(\varrho_1, \mathrm{Res}^G_{G_1} \varrho),$$

of intertwiners between ϱ_1 and the restriction of ϱ to G_1; note that it is non-zero by definition of ϱ_1. We claim that the definition

$$(\varrho_2(g_2)\Phi)(v) = \varrho(1, g_2)\Phi(v)$$

for $g_2 \in G_2$, $\Phi \in E_2$ and $v \in E_1$, defines a representation ϱ_2 of $G_2 = \{1\} \times G_2 \subset G$ on E_2. The point is that because G_1 and G_2, seen as subgroups of G, commute with each other, the k-linear map $\varrho_2(g_2)\Phi$ is still a homomorphism $\varrho_1 \to \varrho$ (and not merely a linear map). Indeed, denoting $\Psi = \varrho_2(g_2)\Phi$, we compute

$$\begin{aligned}\Psi(\varrho_1(g_1)v) &= \varrho(1, g_2)\Phi(\varrho_1(g_1)v)\\ &= \varrho(1, g_2)(\varrho(g_1, 1)\Phi(v)) \qquad (\text{since } \Phi \in E_2)\\ &= \varrho(g_1, 1)(\varrho(1, g_2)\Phi(v)) = \varrho(g_1, 1)\Psi(v),\end{aligned}$$

for all $v \in E_1$, which is to say, $\Psi \in E_2$.

Now we define a k-linear map

$$\Theta \begin{cases} E_1 \otimes E_2 & \longrightarrow & E \\ v \otimes \Phi & \mapsto & \Phi(v), \end{cases}$$

and we claim that Θ is an intertwiner between $\varrho_1 \boxtimes \varrho_2$ and ϱ. Indeed, we can check this on pure tensors: for $g_i \in G_i$, $v \otimes \Phi \in E_1 \otimes E_2$, we have

$$\begin{aligned}\Theta(\varrho_1 \boxtimes \varrho_2(g_1, g_2)(v \otimes \Phi)) &= \Theta(\varrho_1(g_1)v \otimes \varrho_2(g_2)\Phi)\\ &= (\varrho_2(g_2)\Phi)(\varrho_1(g_1)v)\\ &= \varrho(1, g_2)\Phi(\varrho_1(g_1)v)\\ &= \varrho(1, g_2)\varrho(g_1, 1)\Phi(v)\\ &= \varrho(g_1, g_2)\Phi(v) = \varrho(g_1, g_2)\Theta(v \otimes \Phi)\end{aligned}$$

(this must be written down by yourself to not look like gibberish).

Now to show that Θ is bijective, we can use the following trick. Let $F_2 \subset E_2$ be any irreducible subrepresentation (of G_2), the action being denoted ϱ_2'; restricting Θ to $E_1 \otimes F_2$ gives an intertwiner

$$\varrho_1 \boxtimes \varrho_2' \longrightarrow E.$$

By the first part of Proposition 2.3.17, that we proved at the beginning, the representation $\varrho_1 \boxtimes \varrho_2'$ of G is irreducible, and so is E by assumption; moreover, if $v_0 \neq 0$ is a vector in E_1 and $0 \neq \Phi_0 \in F_2$, then $\Theta(v_0 \otimes \Phi_0) = \Phi_0(v_0)$ is non-zero (by Schur's Lemma, because Φ_0 is an embedding of the irreducible representation ϱ_1 in ϱ, it is injective). Thus Θ, restricted to $\varrho_1 \boxtimes \varrho_2'$, is non-zero, and again by Schur's Lemma, must be an isomorphism.

We are in fact done proving that ϱ is an external tensor product, but we will continue with some (minor) additional work that shows that, in fact, Θ itself is bijective. For this, we just need to show that

$$\dim(E_1 \otimes E_2) = \dim E = \dim(E_1)\dim(E_2'),$$

(since we now know that Θ is surjective) or equivalently that $\dim E_2' = \dim E_2$. But

$$E_2 = \operatorname{Hom}_{G_1}(\varrho_1, \operatorname{Res}^G_{G_1} \varrho),$$

and, by fixing a basis (v_j) of the space of ϱ_2', we see that the restriction to G_1 of $\varrho = \varrho_1 \boxtimes \varrho_2'$ is the direct sum

$$\bigoplus_j E_1 \otimes kv_j \simeq (\dim \varrho_2')E_1,$$

so that the dimension of E_2 is given by

$$\dim E_2 = (\dim \varrho_2') \dim \operatorname{Hom}_{G_1}(\varrho_1, \varrho_1) = \dim \varrho_2'$$

by the last part of Schur's Lemma (we use again the fact that k is algebraically closed).

Finally, coming back to the general situation, note that this last observation on the restriction of $\varrho_1 \boxtimes \varrho_2$ to G_1 (and the analogue for G_2) show that ϱ_1 and ϱ_2 are indeed unique up to isomorphism, by the Jordan-Hölder-Noether Theorem. □

Exercise 2.7.29. In this exercise we prove Proposition 2.3.17 when k is not necessarily algebraically closed.

(1) Show that the second part (every finite-dimensional irreducible k-representation of $G = G_1 \times G_2$ is an external tensor product) will follow from the first. [Hint: Follow the same argument as in the algebraically closed case.]

(2) Consider two irreducible finite-dimensional representations ϱ_i of G_i, acting on E_i, and let $\varrho = \varrho_1 \boxtimes \varrho_2$. Show that if $F \neq 0$ is an irreducible subrepresentation of $E_1 \otimes E_2$ (under the action of ϱ) there exists a non-zero intertwiner

$$E_1 \boxtimes F_2 \longrightarrow F$$

for some irreducible representation space F_2 of G_2. [Hint: Use the ideas of the second part in the algebraically closed case.]

(3) Conclude by showing that necessarily $F_2 \simeq E_2$, and $F \simeq E_1 \boxtimes E_2$. [Hint: Show that E_2 is a composition factor of F, and note that F_2 is the only possible composition factor for the restriction of $E_1 \boxtimes F_2$ to G_2.]

2.7.4. Burnside's theorem and its generalizations, 2. As promised, we now explain how to recover the results of the previous section in a style closer (maybe) to that of Frobenius. Even for readers who have understood the arguments already used, this may be useful. In fact, the proofs are simpler, but not so well motivated. The viewpoint is to start this time by determining directly the isotypic components of the regular representation.

Proposition 2.7.30 (Isotypic component of the regular representation). *Let k be an algebraically closed field and G a group. For any finite-dimensional irreducible k-representation*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of G, the ϱ-isotypic component $M(\varrho) \subset C_k(G)$ of the regular representation is isomorphic to a direct sum of $\dim(\varrho)$ *copies of ϱ, and is spanned by the matrix coefficients of ϱ.*

Proof. We start with the realization (2.22) of $C_k(G)$ as the induced representation $\mathrm{Ind}_1^G(\mathbf{1})$, and then apply Frobenius Reciprocity (2.23), which gives us linear isomorphisms

$$\mathrm{Hom}_G(\varrho, C_k(G)) = \mathrm{Hom}_G(\varrho, \mathrm{Ind}_1^G(\mathbf{1})) \simeq \mathrm{Hom}_1(\mathrm{Res}_1^G(\varrho), \mathbf{1}) = \mathrm{Hom}_k(E, k) = E'.$$

Thus the isotypic component $M(\varrho)$, which is equal to the image of the injective G-homomorphism

$$\begin{cases} \mathrm{Hom}_G(\varrho, C_k(G)) \otimes E & \longrightarrow & C_k(G) \\ \Phi \otimes v & \mapsto & \Phi(v) \end{cases}$$

(see Lemma 2.7.18) is isomorphic to $E' \otimes E$ as a representation of G, where E' has the trivial action of G, and E the action under ϱ. This is the same as the direct sum of $\dim(E)$ copies of ϱ.

We now show, also directly, that the image of the linear map above is spanned by matrix coefficients. Indeed, given $\lambda \in E'$, the corresponding homomorphism

$$\Phi_\lambda : E \longrightarrow C_k(G)$$

in $\mathrm{Hom}_G(\varrho, C_k(G))$ is given, according to the recipe in (2.26), by

$$\Phi_\lambda(v) = (x \mapsto \lambda(\varrho(x)v)) = f_{v,\lambda},$$

i.e., its values are indeed matrix coefficients. □

Re-proof of Theorem 2.7.24. First of all, we recover Burnside's irreducibility criterion. Consider an irreducible finite-dimensional representation $\varrho : G \longrightarrow \mathrm{GL}(E)$. We know from the proposition that

$$\dim M(\varrho) = \dim(\varrho)^2.$$

Since $M(\varrho)$ is spanned by the $\dim(\varrho)^2$ matrix coefficients

$$f_{v_i,\lambda_j}, \quad 1 \leqslant i, j \leqslant \dim \varrho$$

associated with any basis (v_i) of E and the dual basis (λ_j) of E', these must be independent. But then if we consider the matrix representation $\varrho^{\boldsymbol{m}}$ of ϱ in the basis (v_i), that means that there is no non-trivial linear relation between the coefficients of the $\varrho^{\boldsymbol{m}}(g)$, and hence – by duality – the span of those matrices must be the space of all matrices of size $\dim \varrho$, which means that the linear span of the $\varrho(g)$ in $\mathrm{End}(E)$ is equal to $\mathrm{End}(E)$. This recovers Burnside's criterion (since the converse direction was immediate).

Now consider finitely many irreducible representations (ϱ_i) which are pairwise non-isomorphic. The subspaces $M(\varrho_i) \subset C_k(G)$ are in direct sum (Lemma 2.7.19: the intersection of any one with the sum of the others is a subrepresentation where no composition factor is permitted, hence it is zero), and this means that the matrix coefficients of the ϱ_i must be linearly independent – in the sense of the statement of Theorem 2.7.24. □

2.7.5. Characters of finite-dimensional representations. As another consequence of Theorem 2.7.24, we see that if we are given one matrix coefficient of each of ϱ_1 and ϱ_2, some irreducible k-representations of G, both finite-dimensional, we are certain that they will be distinct functions if ϱ_1 and ϱ_2 are not isomorphic. The converse is not true, since even a single representation has typically many matrix coefficients. However, one can combine some of them in such a way that one obtains a function which only depends on the representation up to isomorphism, and which *characterizes* (finite-dimensional) irreducible representations, up to isomorphism.

Definition 2.7.31 (Characters). Let G be a group and k a field.

(1) A *character* of G over k, or k-character of G, is any function $\chi : G \longrightarrow k$ of the type

$$\chi(g) = \mathrm{Tr}\, \varrho(g)$$

where ϱ is a finite-dimensional k-representation of G. One also says that χ is the character of G.

(2) An *irreducible character* of G over k is a character associated to an irreducible k-representation of G.

We will typically write χ_ϱ for the character of a given representation ϱ. Then we have:

Corollary 2.7.32. *Let G be a group and let k be an algebraically closed field.*

(1) *Two irreducible finite-dimensional representations ϱ_1 and ϱ_2 of G are isomorphic if and only if their characters are* equal *as functions on G. More generally, the characters of the irreducible finite-dimensional representations of G, up to isomorphism, are linearly independent in $C_k(G)$.*

(2) *If k is of characteristic zero, then two finite-dimensional semisimple representations of G are isomorphic if and only if their characters are equal.*

Proof. If two representations (irreducible or not, but finite-dimensional) ϱ_1 and ϱ_2 are isomorphic, with $\Phi : E_1 \to E_2$ giving this isomorphism, we have

$$\Phi \circ \varrho_1(g) \circ \Phi^{-1} = \varrho_2(g)$$

for all $g \in G$, and hence

$$\mathrm{Tr}(\varrho_1(g)) = \mathrm{Tr}(\varrho_2(g))$$

so that their characters are equal.

To prove (1), we note simply that the character $\mathrm{Tr}\,\varrho(g)$ is a sum of (diagonal) matrix coefficients: if (v_i) is a basis of the space of ϱ, with dual basis (λ_i), we have

$$\chi_\varrho = \sum_i f_{v_i,\lambda_i}$$

(i.e,

$$\mathrm{Tr}\,\varrho(g) = \sum_i \langle \lambda_i, \varrho(g)v_i \rangle,$$

for all $g \in G$). Hence an equality

$$\chi_{\varrho_1} = \chi_{\varrho_2}$$

is a linear relation between certain matrix coefficients of ϱ_1 and ϱ_2 respectively. If ϱ_1 and ϱ_2 are irreducible but not isomorphic, it follows that such a relation is impossible.

Similarly, expanding in terms of matrix coefficients, we see that any linear relation

$$\sum_\pi \alpha(\pi)\chi_\pi = 0$$

among characters of the irreducible finite-dimensional k-representations of G (taken up to isomorphism) must have $\alpha(\pi) = 0$ for all π.

Finally, let ϱ be a semisimple k-representation. If we write the decomposition

$$\varrho \simeq \bigoplus_\pi n_\varrho(\pi)\pi$$

in terms of the irreducible finite-dimensional k-representations π of G (up to isomorphism), with $n_\varrho(\pi) \geqslant 0$ the multiplicity of π in ϱ, we find a corresponding decomposition of the character

$$\chi_\varrho = \sum_\pi n_\varrho(\pi)\chi_\pi.$$

By (1), if $\chi_{\varrho_1} = \chi_{\varrho_2}$ for two (semisimple finite-dimensional) representations, we must have $n_{\varrho_1}(\pi) = n_{\varrho_2}(\pi)$ for all π. *But* this equality is an equality in k (the characters are k-valued functions); if k has characteristic zero, this implies the corresponding equality of integers, from which we see that ϱ_1 and ϱ_2 are indeed isomorphic. But if k has positive characteristic p, it is possible that the integer $n_{\varrho_1}(\pi) - n_{\varrho_2}(\pi)$ be a multiple of p for all π, and then the characters are still the same. □

Example 2.7.33 (A zero character). Consider any subgroup G of the group $U_n(\mathbf{Z}/p\mathbf{Z})$ of upper-triangular $n \times n$-matrices with coefficients in $\mathbf{Z}/p\mathbf{Z}$, such that all diagonal coefficients are 1, e.g.,

$$U_3(\mathbf{Z}/p\mathbf{Z}) = \Big\{ \begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} \mid a,\ b,\ c \in \mathbf{Z}/p\mathbf{Z} \Big\}.$$

With $k = \mathbf{Z}/p\mathbf{Z}$, the inclusion in $\mathrm{GL}_n(k)$ gives a k-representation

$$\varrho : U_n(\mathbf{Z}/p\mathbf{Z}) \longrightarrow \mathrm{GL}(k^n).$$

Then, if n is a multiple of p, we have

$$\chi_\varrho(g) = 0$$

for all g, despite the fact that ϱ is not trivial. (Here, of course, ϱ is not semisimple and has n trivial composition factors, but n is 0 in the field k.)

EXAMPLE 2.7.34 (The character of the regular representation). Let G be a finite group, and k a field (so that the space $C_k(G)$ of the regular representation of G is finite-dimensional). Then its character is given by

$$\chi(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{if } g \neq 1. \end{cases} \tag{2.43}$$

Indeed, we can take as a basis of $C_k(G)$ the family of functions $(\delta_x)_{x\in G}$ equal to 1 at $g = x$ and 0 elsewhere. Then

$$\mathrm{reg}(g)\delta_x = \delta_{xg^{-1}}$$

for all g and $x \in G$. This means that $\mathrm{reg}(g)$ acts on the basis elements by permuting them, so that the corresponding matrix is a permutation matrix. The trace of $\mathrm{reg}(g)$ is therefore the number of fixed points of this permutation, but we see that the formula that $\mathrm{reg}(g)\delta_x = \delta_x$ if and only if $g = 1$, and then x is arbitrary. This gives the formula we claimed. (Note that here also, if the order of the group G is zero in k, the character becomes identically zero.)

Corollary 2.7.32 is quite remarkable. It gives a tool to study representations of a group using only functions on G, and – especially for finite and compact groups – it is so successful that in some cases, one knows all the characters of irreducible representations of a group – as explicit functions – *without* knowing explicit descriptions of the corresponding representations!

Part of the appeal of characters is that they are quite manageable in terms of computation. The following summarizes the simplest such results:

PROPOSITION 2.7.35 (Formalism of characters). *Let G be a group and let k be a field. For any finite-dimensional k-representation ϱ of G, the character χ_ϱ satisfies*

$$\chi_\varrho(gxg^{-1}) = \chi_\varrho(x), \qquad \chi_\varrho(xg) = \chi_\varrho(gx) \tag{2.44}$$

for all g, $x \in G$.

Moreover, we have the identities

$$\begin{aligned} \chi_\varrho(1) &= \dim \varrho \quad \text{(seen as an element of } k\text{)}, \\ \chi_{\varrho_1\oplus\varrho_2} &= \chi_{\varrho_1} + \chi_{\varrho_2}, \\ \chi_{\varrho_1\otimes\varrho_2} &= \chi_{\varrho_1}\chi_{\varrho_2}, \\ \chi_{\tilde{\varrho}}(g) &= \chi_\varrho(g^{-1}). \end{aligned}$$

If ϱ is finite-dimensional and has distinct composition factors ϱ_i, with multiplicities $n_i \geqslant 1$, then

$$\chi_\varrho = \sum_i n_i \chi_{\varrho_i}.$$

Moreover if $H \subset G$ is a subgroup, we have

$$\chi_{\mathrm{Res}_H^G(\varrho)}(h) = \chi_\varrho(h) \textit{ for all } h \in H,$$

and if H is a finite-index subgroup, we have

$$\chi_{\mathrm{Ind}_H^G(\varrho)}(g) = \sum_{\substack{s\in H\backslash G \\ sgs^{-1}\in H}} \chi_\varrho(sgs^{-1}). \tag{2.45}$$

The two statements in (2.44) are equivalent, and state that the value of a character at some $x \in G$ only depends on the conjugacy class of x in G. Functions with this property are usually called *class functions* on G, or *central functions*.

Note that we restrict to a finite-index subgroup for induction because otherwise the dimension of the induced representation is not finite.

The character formula makes sense because the property that sgs^{-1} be in H, and the value of the trace of $\varrho(sgs^{-1})$, are both unchanged if s is replaced by any other element of the coset Hs. It may also be useful to observe that one can *not* use the invariance of ϱ_ϱ under conjugation to remove the s in $\chi_\varrho(sgs^{-1})$, since ϱ is only a representation of H, and $s \notin H$ (except for the coset H itself).

Note also that by taking $g = 1$, it leads – as it should – to the formula

$$\dim \operatorname{Ind}_H^G(\varrho) = [H : G] \dim \varrho$$

of Proposition 2.3.8, if the field k has characteristic zero (in which case the equality in k gives the same in $\mathbf{Z}$).

Proof. The first formulas are direct consequences of the definitions and the properties of the trace of linear maps. Similarly, the formula for the restriction is clear, and only the case of induction requires proof.

Let E be the space on which ϱ acts, and let F be the space

$$F = \{\varphi : G \to E \mid \varphi(hx) = \varrho(h)\varphi(x) \text{ for } h \in H,\ x \in G\}$$

of the induced representation. We will compute the trace by decomposing F (as a linear space) conveniently, much as was done in the proof of Proposition 2.3.8 when computing the dimension of F (which is also, of course, the value of the character at $g = 1$). First of all, for any $s \in G$, let $F_s \subset F$ be the subspace of those $\varphi \in F$ which vanish for all $x \notin Hs$; thus F_s only depends on the coset $Hs \in H\backslash G$. We then have a direct sum decomposition

$$F = \bigoplus_{s \in H\backslash G} F_s,$$

(where the components of a given φ are just obtained by taking the restrictions of φ to the cosets Hs and extending this by zero outside Hs.)

Now, for a fixed $g \in G$, the action of $\operatorname{Ind}(g)$ on F is given by the regular representation

$$\operatorname{Ind}(g)\varphi(x) = \varphi(xg).$$

It follows from this that $\operatorname{Ind}(g)$ *permutes* the subspaces F_s, and more precisely that $\operatorname{Ind}(g)$ sends F_s to $F_{sg^{-1}}$. In other words, in terms of the direct sum decomposition above, the action of $\operatorname{Ind}(g)$ is a “block permutation matrix”. Taking the trace (it helps visualizing this as a permutation matrix), we see that it is the sum of the trace of the maps induced by $\operatorname{Ind}(g)$ over those cosets $s \in H\backslash G$ for which $Hsg^{-1} = Hs$, i.e., over those s for which $sgs^{-1} \in H$.

Now for any $s \in G$ with $sgs^{-1} \in H$, we compute the trace of the linear map

$$\pi_{s,g} : F_s \longrightarrow F_s$$

induced by $\operatorname{Ind}(g)$. To do this, we use the fact – already used in Proposition 2.3.8 – that F_s is isomorphic to E. More precisely, there are reciprocal k-linear isomorphisms

$$E \xrightarrow{\alpha} F_s \xrightarrow{\beta} E$$

such that

$$\beta(\varphi) = \varphi(s)$$

on the one hand, and $\alpha(v)$ is the element of F_s mapping hs to $\varrho(h)v$ (and all $x \notin Hs$ to 0.) The fact that α and β are inverses of each other is left for the reader to (it is contained in the proof of Proposition 2.3.8).

Thus the trace of $\mathrm{Ind}(g)$ is the sum, over those s, of the trace of the linear map on E given by $\beta \circ \pi_{s,g} \circ \alpha$. But – and this shouldn't be much of a surprise – this map is simply given by

$$\varrho(sgs^{-1}) \,:\, E \longrightarrow E,$$

with trace $\chi_\varrho(sgs^{-1})$ (which is defined because $sgs^{-1} \in H$, of course). The stated formula follows by summing over the relevant s.

We check the claim: given $v \in E$, and $\varphi = \alpha(v)$, we have

$$\begin{aligned}(\beta \circ \pi_{s,g} \circ \alpha)(v) &= \mathrm{Ind}(g)\varphi(s) = \varphi(sg)\\ &= \varphi((sgs^{-1})s) = \varrho(sgs^{-1})\varphi(s) = \varrho(sgs^{-1})v,\end{aligned}$$

using the definitions of α and β. □

Example 2.7.36. The formula for an induced character may look strange of complicated at first. In particular, it is probably not clear just by looking at the right-hand side that it *is* the character of a representation of G! However, we will see, here and especially in Chapter 4, that the formula is quite flexible and much nicer than it may seem.

(1) Example 2.7.34 is also a special case of the formula (2.45) for the character of an induced representation. Indeed, we know that the regular representation reg of a group G (over a field k) is the same as the induced representation $\mathrm{Ind}_1^G(\mathbf{1})$ of the trivial representation of the trivial subgroup $\{1\}$ of G. Hence

$$\chi_{\mathrm{reg}}(g) = \sum_{\substack{s \in G\\ sgs^{-1}=1}} 1,$$

which leads to the formula (2.43), since the conjugacy class of 1 is reduced to $\{1\}$.

(2) Generalizing this, let ϱ be the permutation representation (Section 2.6.2) associated with the action of G on a finite set X. Then we have

$$\chi_\varrho(g) = |\{x \in X \mid g \cdot x = x\}|,$$

i.e., the character value at g is the number of fixed points of g acting on X. Indeed, in the basis (e_x) of the space of ϱ, each $\varrho^{\boldsymbol{m}}(g)$ is a permutation matrix, and its trace is the number of non-zero (equal to 1) diagonal entries. These correspond to those $x \in X$ where $\varrho(g)x = e_{gx}$ is equal to e_x, i.e., to the fixed points of g.

(3) Let $H \lhd G$ be a normal subgroup of G of finite index, and ϱ a finite-dimensional representation of H. Let $\pi = \mathrm{Ind}_H^G(\varrho)$. Then

$$\chi_\pi(g) = 0$$

for $g \notin H$, since the condition $sgs^{-1} \in H$ means $g \in s^{-1}Hs = H$. For $h \in H$, on the other hand, we have $shs^{-1} \in H$ for all s, and thus

$$\chi_\pi(h) = \sum_{s \in G/H} \chi_\varrho(shs^{-1}).$$

Exercise 2.7.37. Show directly using the character formula that if H is a subgroup of finite index in G, the characters on both sides of the projection formula

$$\mathrm{Ind}_H^G(\varrho_2 \otimes \mathrm{Res}_H^G(\varrho_1)) \simeq \mathrm{Ind}_H^G(\varrho_2) \otimes \varrho_1$$

are identical functions on G.

Example 2.7.38 (Characters of $\mathrm{SL}_2(\mathbf{C})$). Consider $G = \mathrm{SL}_2(\mathbf{C})$ and the representations V_m defined in Example 2.6.1 for $m \geqslant 0$. We can compute the character of V_m, to some extent, by using the basis of monomials $e_i = X^i Y^{m-i}$ of the space V_m: by definition, if

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

we have

$$\begin{aligned}
\varrho_m(g)(e_i) &= (aX + cY)^i (bX + dY)^{m-i} \\
&= \sum_{k=0}^{i} \sum_{l=0}^{m-i} \binom{i}{k} \binom{m-i}{l} a^k c^{i-k} b^l d^{m-i-l} X^{k+l} Y^{m-k-l} \\
&= \sum_{j=0}^{m} \Bigl(\sum_{k+l=j} \binom{i}{k} \binom{m-i}{l} a^k c^{i-k} b^l d^{m-i-l} \Bigr) e_j
\end{aligned}$$

by binomial expansion. The diagonal coefficient here is

$$\sum_{k+l=i} \binom{i}{k} \binom{m-i}{l} a^k c^{i-k} b^l d^{m-i-l},$$

and hence

$$\chi_{\varrho_m}(g) = \sum_{i=0}^{m} \sum_{k+l=i} \binom{i}{k} \binom{m-i}{l} a^k c^{i-k} b^l d^{m-i-l}.$$

This may – or not – look forbidding. However, if one remembers that the value of the character at g depends only on the conjugacy class of g, one can simplify this, at least for certain elements. Suppose for instance that g is diagonalizable, hence is conjugate to some matrix

$$t_\lambda = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$$

for some $\lambda \in \mathbf{C}^\times$ (this will be true very often, e.g., whenever the eigenvalues of g are distinct). The computation of $\chi_{\varrho_m}(g)$ is then much easier: we have indeed

$$\varrho_m(t_\lambda) e_i = (\lambda X)^j (\lambda^{-1} Y)^{m-i} = \lambda^{2i-m} e_i$$

for $0 \leqslant i \leqslant m$. Hence, we obtain the formula

$$\chi_{\varrho_m}(g) = \chi_{\varrho_m}(t_\lambda) = \lambda^{-m} + \lambda^{-m+2} + \cdots + \lambda^{m-2} + \lambda^m = \frac{\lambda^{m+1} - \lambda^{-m-1}}{\lambda - \lambda^{-1}}.$$

(This computation corresponds to the fact, already seen in Example 2.7.9, that the restriction of V_m to the diagonal subgroup T is the direct sum (2.37).

If we specialize even further to $\lambda = e^{i\theta}$ with $\theta \in \mathbf{R}$ (i.e., to t_λ being a unitary matrix) we obtain

$$\chi_{\varrho_m}\Bigl(\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix}\Bigr) = \frac{\sin((m+1)\theta)}{\sin(\theta)}, \tag{2.46}$$

(with $\theta = 0$ mapping of course to $m + 1$); these character values are simple – and fundamental! – trigonometric functions.

Suppose on the other hand that $g = u_t$ is conjugate to an upper-triangular matrix

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$$

with $t \in \mathbf{C}$. Then we have

$$\varrho_m(u_t)e_i = X^i(tX+Y)^{m-i}$$

and if we expand the second term, we see quickly that this is of the form

$$X^iY^{m-i} + (\text{combination of } e_j \text{ with } j > i),$$

leading in particular to

$$\chi_{\varrho_m}(u_t) = m+1$$

for all t.

Exercise 2.7.39. Prove that if $m \geqslant n \geqslant 0$, the characters of the two sides of the Clebsch-Gordan formula (2.31) coincide for as large a set of (conjugacy classes of) $g \in \mathrm{SL}_2(\mathbf{C})$ as you can.

The first part of Proposition 2.7.35 shows that the subset of $C_k(G)$ whose elements are characters of finite-dimensional k-representations of G is stable under addition and multiplication. It is therefore quite natural to consider the abelian group generated by all characters, as a subgroup of the additive group of $C_k(G)$. Indeed, the tensor product formula shows that this group is in fact a ring.

Definition 2.7.40 (Generalized, or virtual, characters). Let G be a group and let k be a field of characteristic zero. The *character ring* $R_k(G)$ of *generalized characters* of G over k, is the ring generated, as an abelian group, by the characters of finite-dimensional k-representations of G. The elements of $R_k(G)$ are also called *virtual characters*. The *dimension* of a virtual character χ is defined to be $\chi(1) \in \mathbf{Z} \subset k$.

Note that $R_k(G)$ is *not* a k-vector space: we do not allow linear combinations of characters with coefficients which are not integers. Concretely, a virtual character $\chi \in R_k(G)$ is a function

$$\chi : G \to k$$

of the form

$$\chi = \chi_{\varrho_1} - \chi_{\varrho_2}$$

for some finite-dimensional k-representations ϱ_1 and ϱ_2 of G; note that the latter are by no means unique.

We will present some applications of the character ring in Section 4.8.1.

Example 2.7.41 (Induced representations as ideal in $R_k(G)$). Consider the ring $R_k(G)$ of a group G and a homomorphism $\phi : H \to G$ with image $\mathrm{Im}(\phi)$ of finite index in G. Now consider the subgroup I_ϕ of $R_k(G)$ generated – as abelian group – by all characters of induced representations $\mathrm{Ind}_H^G(\varrho)$, where ϱ is a finite-dimensional representation of H. Then the projection formula (Proposition 2.3.12) shows that $I_\phi \subset R_k(G)$ is *an ideal*, i.e., $\chi_1\chi_2 \in I_\phi$ if χ_1 in I_ϕ and χ_2 is arbitrary.

We will say (much) more about characters, for finite and compact groups, in Chapters 4 and 5.

2.8. Conclusion

We have now built, in some sense, the basic foundations of representation theory, trying to work in the greatest generality. Interestingly, there are some results of comparable generality which the techniques we have used (which involve basic abstract algebra) are – as far as we know! – unable to prove. These require the theory of *algebraic groups*. A very short introduction, with a discussion of some of the results it leads to, can be found in the beginning of Chapter 7.

Before closing this chapter, we can use the vocabulary we have built to ask: What are the fundamental questions of representation theory? The following are maybe the two most obvious ones:

- (Classification) For a group G, one may want to classify all its irreducible representations (say over the complex numbers), or all representations of a certain type. This is possible in a number of very important cases, and is often of tremendous importance for applications; one should think here of a group G which is fairly well-understood from a group-theoretic point of view.
- (Decomposition) Given a group G again, and a natural representation E of G (say over $\mathbf{C}$), one may ask to find explicitly the irreducible components of E, either as summands if E is semisimple, or simply as Jordan-Hölder factors.

Both are crucial to the so-called "Langlands Program" in modern number theory.

CHAPTER 3

Variants

We discuss here some of the variants of the definition of representations that we have used. Many of them are very important topics in their own right, but we will only come back, usually rather briefly, to some of them.

3.1. Representations of algebras

If, instead of a group G, we consider an *algebra* A over a field k, i.e., a ring which has a compatible structure of k-vector space, the analogue of a representation of A is an algebra homomorphism

$$A \xrightarrow{\varrho} \mathrm{End}_k(E).$$

This may be called a *representation of* A, but it is more usual to focus on E and to note that such a map defines a structure of A-module on the vector space E by

$$a \cdot v = \varrho(a)v$$

for $a \in A$ and $v \in E$.

It is important for certain aspects of representation theory that the k-representations of a group G can be understood in this language. Associated to G is the vector space $k(G)$ freely generated by G (which already briefly appeared at the beginning of the previous chapter); this has in fact the structure of a k-algebra if one defines the product on $k(G)$ by

$$[g] \cdot [h] = [gh],$$

and expand it by linearity, where we use temporarily $[g]$ to indicate the g-th basis vector of $k(G)$. Thus any elements a, $b \in k(G)$ can be written as

$$a = \sum_{g \in G} \lambda_g [g], \qquad b = \sum_{h \in G} \mu_h [h]$$

respectively (with only finitely many non-zero coefficients, if G is infinite), and their product is

$$ab = \sum\sum_{g,h \in G} \lambda_g \mu_h [g][h] = \sum_{x \in G} \Bigl(\sum_{gh=x} \lambda_g \mu_h \Bigr) [x].$$

Note that if G is not abelian, this algebra is not commutative.

Example 3.1.1. Let G be a finite group, k a field, and consider the element

$$s = \sum_{g \in G} [g] \in k(G).$$

Then by expanding the square, we find in $k(G)$ that

$$s^2 = \sum\sum_{g,h \in G} [gh] = \sum_{x \in G} \Bigl(\sum_{[g][h]=x} 1 \Bigr) x = |G| s. \tag{3.1}$$

Similarly, multiplying on the left or the right with $x \in k(G)$, we get

$$s[x] = \sum_{g \in G} [g][x] = \sum_{g \in G} [gx] = \sum_{h \in G} [h] = s, \qquad [x]s = s. \tag{3.2}$$

By linearity, we see that $sa = as$ for all $a \in k(G)$; thus the element s is in the *center* of the algebra $k(G)$.

Definition 3.1.2 (Group algebra). Let k be a field and G a group. The algebra $k(G)$ defined above is called the *group algebra* of G over k.

The characteristic algebraic property of the group algebra is the following:

Proposition 3.1.3 (Representations as modules over the group algebra). *Let G be a group and k a field. If*

$$G \xrightarrow{\varrho} \mathrm{GL}(E)$$

is a k-representation of E, then E has a structure of $k(G)$-module defined by extending ϱ by linearity, i.e.,

$$\Bigl(\sum_{g \in G} \lambda_g [g]\Bigr) \cdot v = \sum_{g \in G} \lambda_g \varrho(g) v.$$

Conversely, any $k(G)$-module E inherits a representation of G by restricting the "multiplication by a" maps to the basis vectors $[g]$, i.e., by putting

$$\varrho(g) = (v \mapsto [g] \cdot v).$$

Sometimes, to avoid dropping the reference to the representation ϱ, one writes

$$\varrho(a) v = \sum_{g \in G} \lambda_g \varrho(g) v$$

for an element $a \in k(G)$ as above. Thus $\varrho(a)$ becomes an element in $\mathrm{Hom}_k(E, E)$.

Here are some reasons why this approach can be very useful:

- It gives access to all the terminology and results of the theory of algebras; in particular, for any ring A, the notions of sums, intersections, direct sums, etc, of A-modules, are well-defined and well-known. For $A = k(G)$, they correspond to the definitions already given in the previous chapter for linear representations. More sophisticated constructions are however more natural in the context of algebras. For instance, one may consider the ideals of $k(G)$ and their properties, which is not something so natural at the level of the representations themselves.
- The $k(G)$-modules parametrize, in the almost tautological way we have described, the k-representations of G. It may happen that one wishes to concentrate attention on a special class $\mathcal{C}$ of representations, characterized by some property. It happens, but very rarely, that these representations correspond in a natural way to all (or some of) the representations of another group $G_{\mathcal{C}}$ (an example is to consider for $\mathcal{C}$ the class of one-dimensional representations, which correspond to those of the derived group $G/[G, G]$, as in Proposition 2.6.6), but it may happen more often that there is a natural k-algebra $A_{\mathcal{C}}$ such that *its* modules correspond precisely (and "naturally") to the representations in $\mathcal{C}$.

Partly for reasons of personal habit (or taste), we won't exploit the group algebra systematically in this book. This can be justified by the fact that it is not absolutely necessary at the level we work. But we will say a bit more, e.g., in Section 4.3.6, and Exercise 4.3.29 describes a property of representations with cyclic vectors which is much more natural from the point of view of the group algebra.

Exercise 3.1.4 (Representations with a fixed vector under a subroup). Let G be a finite group. Consider a subgroup $H \subset G$; note that $\mathbf{C}(H)$ is naturally a subring of $\mathbf{C}(G)$.

(1) Show that

$$\mathcal{H} = \{a \in k(G) \mid hah' = a \text{ for all } h,\ h' \in H\}$$

is a subalgebra of $k(G)$, and that it is generated as k-vector space by the characteristic functions of all double classes

$$HxH = \{hxh' \mid h,\ h' \in H\} \subset G.$$

(2) Show that if $\varrho : G \longrightarrow \mathrm{GL}(E)$ is a representation of G, the subspace E^H is stable under multiplication by $\mathcal{H}$, i.e., that E^H is an $\mathcal{H}$-module.

(3) Show that if $\Phi \in \mathrm{Hom}_G(E, F)$, the restriction of Φ to E^H is an $\mathcal{H}$-linear map from E^H to F^H.

(4) Let $\varrho : G \longrightarrow \mathrm{GL}(E)$ be a semisimple representation of G such that $E^H \neq 0$. Show that ϱ is irreducible if and only if E^H is simple as an $\mathcal{H}$-module, i.e., if and only if E^H contains no proper non-zero submodule.

(5) Show that if ϱ_1, ϱ_2 are irreducible representations of G on E_1, E_2 respectively, with $E_i^H \neq 0$, the $\mathcal{H}$-modules E_1^H and E_2^H are isomorphic if and only if $\pi_1 \simeq \pi_2$.

We will use the $k(G)$-module structure corresponding to representations a number of times in the remainder of the book. Usually, the notation $[g]$ will be abandoned in doing so, and we will write, e.g.

$$s = \sum_{g \in G} g \in k(G)$$

for the element of Example 3.1.1.

Exercise 3.1.5 (The group ring as "universal" endomorphisms). A fixed $a \in k(G)$ has the feature that, for any representation $\varrho : G \longrightarrow \mathrm{GL}(E)$, there is a corresponding k-linear endomorphism given by its action on E, i.e., the linear map

$$\varepsilon_\varrho^{(a)} \begin{cases} E \longrightarrow E \\ v \mapsto \varrho(a)v \end{cases}$$

in $\mathrm{Hom}_k(E, E)$.

These maps are "functorial" in ϱ, in the sense that for any other representation $\pi : G \longrightarrow \mathrm{GL}(F)$ and any morphism of representations $\Phi \in \mathrm{Hom}_G(\varrho, \pi)$, the square diagram

$$\begin{tikzcd} E \arrow[r, "\varepsilon_\varrho^{(a)}"] \arrow[d, "\Phi"'] & E \arrow[d, "\Phi"] \\ F \arrow[r, "\varepsilon_\pi^{(a)}"] & F \end{tikzcd}$$

is commutative. This exercise shows that this property characterizes the group algebra. Namely consider now any map

$$\varrho \mapsto \varepsilon_\varrho$$

sending a k-representation ϱ of G acting on E to a linear map $\varepsilon_\varrho \in \mathrm{Hom}_k(E, E)$, for which the rules above are valid. We will show that, for some *fixed* $a \in k(G)$, we have $\varepsilon_\varrho = \varepsilon_\varrho^{(a)}$ for all representations ϱ.

(1) Show that there exists $a \in k(G)$ such that $\varepsilon_{k(G)}$ is the linear map

$$\varepsilon_{k(G)}^{(a)} : x \mapsto ax$$

on $k(G)$, seen as a $k(G)$-module by multiplication on the left. [Hint: Consider the maps $\Phi : x \mapsto xb$ on $k(G)$.]

(2) With a as in (1), show that $\varepsilon_\varrho = \varepsilon_\varrho^{(a)}$ for any representation ϱ with a cyclic vector v_0.

(3) Conclude that $\varepsilon_\varrho = \varepsilon_\varrho^{(a)}$ for all representations.

(4) Show that the "universal endomorphisms" associated to a and $b \in k(G)$ are the same if and only if $a = b$.

(5) Show that $a \in k(G)$ is such that $\varepsilon_\varrho^{(a)}$ is in the subspace $\mathrm{End}_G(\varrho)$ of self-intertwiners of ϱ, for all representations ϱ of G, if and only if a is in the center of the group algebra, i.e., if and only if $ax = xa$ for all $x \in k(G)$.

A motivating application of this exercise appears in Section 4.3.6.

3.2. Topological groups

In many applications, it is particularly important to restrict the representations to respect some additional structure. Among these, topological conditions are the most common.

The corresponding structure is that of a *topological group*, i.e., a group G equipped with a topology such that the product map

$$G \times G \to G$$

and the inverse map

$$G \to G$$

are both continuous (with $G \times G$ being given the product topology). There are many examples; for instance, any finite group can be seen as a topological group with the discrete topology; the additive group $\mathbf{R}$ or the multiplicative group $\mathbf{R}^\times$, with the usual euclidean topology, are also topological groups; similarly, for any $n \geqslant 1$, $\mathrm{GL}_n(\mathbf{C})$ is a group with the topology coming from its inclusion in $\mathrm{M}_n(\mathbf{C}) \simeq \mathbf{C}^{n^2}$; and moreover, any subgroup H of a topological group G which is *closed* in G inherits from G a topology and is then a topological group. This includes, for instance, groups like $\mathbf{Z} \subset \mathbf{R}$, $\mathrm{SL}_n(\mathbf{R}) \subset \mathrm{GL}_n(\mathbf{R})$ or $\mathrm{SL}_n(\mathbf{Z}) \subset \mathrm{SL}_n(\mathbf{R})$.

When dealing with a topological group, one typically wishes to restrict the representations which are considered to include some continuity property. This usually means taking the base field to be either $k = \mathbf{C}$ or $\mathbf{R}$, and the vector space to carry a suitable topology. We will restrict our attention to Banach spaces, i.e., k-vector spaces E with a norm $\|\cdot\|$ on E such that E is complete for this norm. Of special interest, in fact, will be Hilbert spaces, when the norm derives from an inner product $\langle\cdot,\cdot\rangle$. As a special case, it is important to recall that if E is a finite-dimensional real or complex vector space, it carries a unique topology of Banach space, which can be defined using an inner product if desired. (Though, as is well-known, there are many equivalent norms which can be used to define this topology.)

Given a topological group G and a Banach space E, the first restriction concerning representations ϱ of G on E is that the operators $\varrho(g)$, $g \in G$, should be *continuous.* Precisely, we denote by $\mathrm{L}(E)$ the space of continuous linear maps

$$T : E \longrightarrow E,$$

and by $\mathrm{BGL}(E)$ the group of those $T \in \mathrm{L}(E)$ which are invertible, i.e., bijective with a continuous inverse. In fact, this last continuity condition is automatic when E is a

Banach space, by the Banach isomorphism theorem. Of course, if $\dim(E) < +\infty$, we have $\mathrm{BGL}(E) = \mathrm{GL}(E)$.

If $\mathrm{L}(E)$ is given any topology, the group $\mathrm{BGL}(E)$ inherits, as a subset of $\mathrm{L}(E)$ a topology from the latter. It is then natural to think of considering representations

$$\varrho : G \longrightarrow \mathrm{BGL}(E)$$

which are continuous. However, as readers familiar with functional analysis will already know, quite a few different topologies are commonly encountered on $\mathrm{L}(E)$. The most natural-looking[1] is the *norm topology*, associated to the norm

$$\|T\|_{\mathrm{L}(E)|} = \sup_{\substack{v\in E \\ \|v\|\leqslant 1}} \|T(v)\| = \sup_{\substack{v\in E \\ v\neq 0}} \frac{\|Tv\|}{\|v\|} = \sup_{\substack{v\in E \\ \|v\|=1}} \|T(v)\| \tag{3.3}$$

defined for $T \in \mathrm{L}(E)$ (in other words, it is a topology of uniform convergence on bounded subsets of E; the equality of those three quantities follows from linearity). But it turns out that asking for homomorphisms to $\mathrm{BGL}(E)$ which are continuous for this topology does not lead to a good theory: there are "too few" representations in that case, as we will illustrate below. The "correct" definition is the following:

Definition 3.2.1 (Continuous representation). Let G be a topological group and $k = \mathbf{R}$ or $\mathbf{C}$. Let E be a Banach space.

(1) A *continuous representation*, often simply called a representation, of G on E is a homomorphism

$$\varrho : G \longrightarrow \mathrm{BGL}(E)$$

such that the corresponding action map

$$\begin{cases} G \times E \longrightarrow E \\ (g, v) \mapsto \varrho(g)v \end{cases}$$

is continuous, where $G \times E$ has the product topology.

(2) If ϱ_1 and ϱ_2 are continuous representations of G, acting on E_1 and E_2 respectively, a homomorphism

$$\varrho_1 \xrightarrow{\Phi} \varrho_2$$

is a *continuous* linear map $E_1 \longrightarrow E_2$ such that

$$\Phi(\varrho_1(g)v) = \varrho_2(g)\Phi(v)$$

for all $g \in G$ and $v \in E_1$.

If $\dim(E) < +\infty$, note that this is indeed equivalent to asking that ϱ be continuous, where $\mathrm{BGL}(E)$ as the topology from its isomorphism with $\mathrm{GL}_n(k)$, $n = \dim(E)$.

Example 3.2.2 (Too many representations). Consider $G = \mathbf{R}$ and $k = \mathbf{C}$. If we consider simply one-dimensional representations

$$\mathbf{R} \to \mathbf{C}^\times,$$

with no continuity assumptions at all, there are "too many" for most purposes. Indeed, as an abelian group, $\mathbf{R}$ is torsion-free, and hence is a free abelian group: selecting a basis[2] $(v_i)_{i\in I}$ of $\mathbf{R}$ as abelian group, we obtain zillions of 1-dimensional representations by deciding simply that

$$v_i \mapsto z_i$$

[1] It is natural, for instance, because $\mathrm{L}(E)$ becomes a Banach space for this norm.

[2] Necessarily uncountable, and not measurable as a subset of $\mathbf{R}$.

for some arbitrary $z_i \in \mathbf{C}$ and extending these by linearity to $\mathbf{R}$, which is the free group generated by the (v_i).

But as soon as we impose some regularity on the homomorphisms $\mathbf{R} \to \mathbf{C}^\times$, we obtain a much better understanding:

Proposition 3.2.3 (Continuous characters of $\mathbf{R}$). *Let $\chi : \mathbf{R} \to \mathbf{C}^\times$ be a* continuous *group homomorphism. Then there exists a unique $s \in \mathbf{C}$ such that*

$$\chi(x) = e^{sx}$$

for all $x \in \mathbf{R}$.

In fact, one can show that it is enough to ask that the homomorphism be *measurable.* The intuitive meaning of this is that, if one can "write down" a homomorphism $\mathbf{R} \to \mathbf{C}^\times$ in any concrete way, or using standard limiting processes, it will automatically be continuous, and hence be one of the ones above.

The proof we give uses differential calculus, but one can give completely elementary arguments (as in [**1**, Example A.2.5]).

Proof. If χ is differentiable, and not merely continuous, this can be done very quickly using differential equations: we have

$$\chi'(x) = \lim_{h\to 0} \frac{\chi(x+h)-\chi(x)}{h} = \chi(x) \lim_{h\to 0} \frac{\chi(h)-1}{h} = \chi(x)\chi'(0)$$

for all $x \in \mathbf{R}$. Denoting $s = \chi'(0)$, any solution of the differential equation $y' = sy$ is given by

$$y(x) = \alpha e^{sx}$$

for some parameter $\alpha \in \mathbf{C}$. In our case, putting $x = 0$ leads to $1 = \chi(0) = \alpha$, and hence we get the result.

We now use a trick to show that any continuous homomorphism $\chi : \mathbf{R} \longrightarrow \mathbf{C}^\times$ is in fact differentiable. We define the primitive

$$\Psi(x) = \int_0^x \chi(u)du$$

(note that this is $s^{-1}(e^{sx}-1)$ if $\chi(x) = e^{sx}$; this formula explains the manipulations to come), which is a differentiable function on $\mathbf{R}$. Then we write

$$\begin{aligned}\Psi(x+t) = \int_0^{x+t} \chi(u)du = \int_0^x \chi(u)du + \int_x^{x+t} \chi(u)du \\ = \Psi(x) + \int_0^t \chi(x+u)du = \Psi(x) + \chi(x)\Psi(t)\end{aligned}$$

for all real numbers x and t. The function Ψ can not be identically zero (its derivative χ would then also be zero, which is not the case), so picking a fixed $t_0 \in \mathbf{R}$ with $\Psi(t_0) \neq 0$, we obtain

$$\chi(x) = \frac{\Psi(x+t_0)-\Psi(x)}{\Psi(t_0)},$$

which is a differentiable function! Thus our first argument shows that $\chi(x) = e^{sx}$ for all x, with $s = \chi'(0)$. □

Example 3.2.4 (Too few representations). Let G be the compact group $\mathbf{R}/\mathbf{Z}$. We now show that, if one insisted on considering as representations only homomorphisms

$$\varrho\,:\,\mathbf{R}/\mathbf{Z}\longrightarrow \mathrm{BGL}(E)$$

which are continuous with respect to the norm topology on $\mathrm{BGL}(E)$, there would be "too few" (similar examples hold for many other groups). In particular, there would be no analogue of the regular representation. Indeed, if f is a complex-valued function on $\mathbf{R}/\mathbf{Z}$ and $t\in\mathbf{R}/\mathbf{Z}$, it is natural to try to define the latter with

$$\varrho(t)f(x)=f(x+t).$$

To have a Banach space of functions, we must impose some regularity condition. Although the most natural spaces turn out to be the L^p spaces, with respect to Lebesgue measure, we start with $E=C(\mathbf{R}/\mathbf{Z})$, the space of continuous functions $f\,:\,\mathbf{R}/\mathbf{Z}\rightarrow\mathbf{C}$. If we use the supremum norm

$$\|f\|=\sup_{t\in\mathbf{R}/\mathbf{Z}}|f(t)|,$$

this is a Banach space. Moreover, the definition above clearly maps a function $f\in E$ to $\varrho(t)f\in E$; indeed, we have

$$\|\varrho(t)f\|=\|f\|$$

(since the graph of $\varrho(t)f$, seen as a periodic function on $\mathbf{R}$, is just obtained from that of f by translating it to the left by t units), and this further says that $\varrho(t)$ is a continuous linear map on E. Hence, ϱ certainly defines a homomorphism

$$\varrho\,:\,\mathbf{R}/\mathbf{Z}\longrightarrow \mathrm{BGL}(E).$$

But now we claim that ϱ is *not* continuous for the topology on $\mathrm{BGL}(E)$ coming from the norm (3.3). This is quite easy to see: in fact, for any $t\neq 0$ with $0<t<1/2$, we have

$$\|\varrho(t)-\mathrm{Id}_E\|_{\mathrm{L}(E)}=\|\varrho(t)-\varrho(0)\|_{\mathrm{L}(E)}\geqslant 1$$

which shows that ϱ is very far from being continuous at 0. To see this, we take as "test" function any $f_t\in E$ which is zero outside $[0,t]$, non-negative, always $\leqslant 1$, and equal to 1 at $t/2$, for instance (where we view $\mathbf{R}/\mathbf{Z}$ as obtained from $[0,1]$ by identifying the end points). Then $\|f_t\|=1$, and therefore

$$\|\varrho(t)-\mathrm{Id}\|_{\mathrm{L}(E)}\geqslant\|\varrho(t)f_t-f_t\|=\sup_{x\in\mathbf{R}/\mathbf{Z}}|f_t(t+x)-f_t(x)|=1$$

since $\varrho(t)f_t$ is supported on the image modulo $\mathbf{Z}$ of the interval $[1-t,t]$ which is *disjoint* from $[0,t]$, apart from the common endpoint t.

However, the point of this is that we had to use a different test function for each t, and for a fixed f, the problem disappears: by uniform continuity of $f\in E$, we have

$$\lim_{t\rightarrow 0}\|\varrho(t)f-f\|=0$$

for any fixed $f\in E$.

Exercise 3.2.5. Show that ϱ defined above on $E=C(\mathbf{R}/\mathbf{Z})$ is a continuous representation in the sense of Definition 3.2.1.

The general formalism of representation theory can, to a large extent, be transferred or adapted to the setting of representations of topological groups. In particular, for finite-dimensional representations, since all operations considered are "obviously continuous", every construction goes through. This applies to direct sums, tensor products, the contragredient, symmetric and exterior powers, subrepresentations and quotients, etc.

Some care is required when considering infinite-dimensional representations. For instance, the definition of subrepresentations and quotient representations, as well as that of irreducible representation, should be adjusted to take the topology into account:

Definition 3.2.6 (Subrepresentations and irreducibility for topological groups). Let G be a topological group, and let

$$\varrho\,:\,G\longrightarrow \mathrm{BGL}(E)$$

be a representation of G on a Banach space E.

(1) A representation π of G is a subrepresentation of ϱ if π acts on a *closed* subspace $F\subset E$ which is invariant under ϱ, and $\pi(g)$ is the restriction of $\varrho(g)$ to F. Given such a subrepresentation π, the quotient ϱ/π is the induced representation on the quotient Banach space E/F with the norm

$$\|v\|_{E/F}=\min\{\|w\|_E\ \mid\ w\,(\mathrm{mod}\,F)=v\}.$$

(2) The representation ϱ is irreducible (sometimes called *topologically irreducible*) if E is non-zero and E has no non-zero proper subrepresentation, i.e., there is no non-zero proper closed subspace F of E which is stable under all $\varrho(g)$, $g\in G$.

If $\dim(E)<+\infty$, since any subspace of F is closed, there is no difference between these definitions and the previous one. But in general the distinction between closed subspaces and general subspaces is necessary to obtain a good formalism. We will see (in examples like that of $\mathrm{SL}_2(\mathbf{R})$) that there do exist infinite-dimensional representations which are topologically irreducible but have a lot of non-zero stable subspaces. These are necessarily dense in the space E, but they may well be distinct.

The second example we give of adapting the formalism is that of the contragredient representation. Given a Banach space E, the dual Banach space is the vector space E' of continuous linear maps $E\longrightarrow\mathbf{C}$ with the norm

$$\|\lambda\|_{E'}=\sup\{|\lambda(v)|\ \mid\ \|v\|_E\leqslant 1\}.$$

Given a representation

$$\varrho\,:\,G\longrightarrow BGL(E)$$

of a topological group G on E, the contragredient $\tilde{\varrho}$ acts on E' by the usual formula

$$\langle g\cdot\lambda,v\rangle=\langle\lambda,\varrho(g^{-1})v\rangle.$$

If E is finite-dimensional, this is obviously continuous. Otherwise, the following gives a simple condition under which the contragredient is continuous:

Lemma 3.2.7 (Continuity of the contragredient). *Let G be a topological group and ϱ a representation of G on the Banach space E, such that*

$$\sup_{g\in G}\|\varrho(g)\|_{\mathrm{L}(E)}<+\infty.$$

Then the contragredient representation on E' is a continuous representation.

Proof. It is enough to check continuity of the action map at the pair $(1,0)\in G\times E'$ (we leave this reduction as an exercise). But we have

$$\|\tilde{\varrho}(g)\lambda\|_{E'}=\sup_{\|v\|\leqslant 1}|\lambda(\varrho(g^{-1})v)|\leqslant M\|\lambda\|$$

with $M=\sup\|\varrho(g)\|$. Thus if λ is close enough to 0 (and g arbitrary), so is $\tilde{\varrho}(g)\lambda$, which means that $\tilde{g}$ is continuous. □

3.3. Unitary representations

When the representation space (for a topological group) is a Hilbert space H, with an inner product[3] $\langle\cdot,\cdot\rangle$, it is natural to consider particularly closely the *unitary representations*, which are those for which the operators $\varrho(g)$ are *unitary*, i.e., preserve the inner product:

$$\langle\varrho(g)v,\varrho(g)w\rangle = \langle v,w\rangle$$

for all $g \in G$, v and $w \in H$.

We present here most basic facts about such representations. They will be considered in more detail first for finite groups, and then – more sketchily – for compact and locally compact groups. For additional information on the general theory of unitary representations, especially with respect to infinite-dimensional cases, we recommend the Appendices to [**1**].

Definition 3.3.1 (Unitary representations, unitarizable representations). Let G be a topological group, which can be an arbitrary group with the discrete topology.

(1) A *unitary representation* of G is a continuous representation of G on a Hilbert space H where $\varrho(g) \in \mathrm{U}(H)$ for all $g \in G$, where $\mathrm{U}(H)$ is the group of unitary operators of G. A morphism $\varrho_1 \to \varrho_2$ of unitary representations, acting on H_1 and H_2 respectively, is a morphism of continuous representations $\varrho_1 \xrightarrow{\Phi} \varrho_2$ (one does *not* require that Φ preserve the inner product.)

(2) An arbitrary representation

$$G \longrightarrow \mathrm{GL}(E)$$

of G on a complex vector space E is *unitarizable* if there exists an inner product $\langle\cdot,\cdot\rangle$ on E, defining a structure of Hilbert space, such that the values of ϱ are unitary for this inner product, and the resulting map $G \longrightarrow \mathrm{U}(E)$ is a unitary representation.

Remark 3.3.2. If H is finite-dimensional and G carries the discrete topology (e.g., if G is finite) it is enough to construct an inner product on the vector space E such that ϱ takes value in $\mathrm{U}(E)$ in order to check that a representation is unitarizable (the continuity is automatic).

The continuity requirement for unitary representations can be rephrased in a way which is easier to check:

Proposition 3.3.3 (Strong continuity criterion for unitary representations). *Let $\varrho : G \longrightarrow \mathrm{U}(H)$ be a homomorphism of a topological group G to a unitary group of some Hilbert space. Then ϱ is a unitary representation, i.e., the action map is continuous, if and only if ϱ is* strongly continuous*: for any fixed $v \in H$, the map*

$$\begin{cases} G \longrightarrow H \\ g \mapsto \varrho(g)v \end{cases}$$

is continuous on G. Equivalently, this holds when these maps are continuous at $g = 1$.

Proof. The joint continuity of the two-variable action map implies that of the maps above, which are one-variable specializations. For the converse, we use the unitarity and a splitting of epsilons...

[3] Recall from the introduction that our inner products are linear in the first variable, and conjugate-linear in the second: $\langle v,\lambda w\rangle = \bar{\lambda}\langle v,w\rangle$ for $\lambda \in \mathbf{C}$.

Let (g, v) be given in $G \times H$. For any $(h, w) \in G \times H$, we have

$$\|\varrho(g)v - \varrho(h)w\| = \|\varrho(h)(\varrho(h^{-1}g)v - w)\| = \|\varrho(h^{-1}g)v - w\|$$

by unitarity. Then, by the triangle inequality, we get

$$\|\varrho(g)v - \varrho(h)w\| \leqslant \|\varrho(h^{-1}g)v - v\| + \|v - w\|.$$

Under the assumption of the continuity of $x \mapsto \varrho(x)v$ when $x \to 1$, this shows that when w is close to v and h close to g, the difference becomes arbitrarily small, which is the continuity of the action map at (g, v). To be precise: given $\varepsilon > 0$, we can first find an open neighborhood U_v of v in H such that $\|v - w\| < \varepsilon$ for $w \in U_v$, and we can use the continuity assumption to find an open neighborhood U_1 of $1 \in G$ such that $\|\varrho(x)v - v\| < \varepsilon$ for $x \in U_1$. Then, when $(h, w) \in gU_1^{-1} \times U_v$, we have

$$\|\varrho(g)v - \varrho(h)w\| \leqslant 2\varepsilon.$$

□

Remark 3.3.4 (The strong topology). The name “strong continuity” refers to the fact that the condition above is equivalent with the assertion that the homomorphism

$$\varrho : G \longrightarrow \mathrm{U}(H)$$

is continuous, with respect to the topology induced on $\mathrm{U}(H)$ by the so-called *strong operator topology* on $\mathrm{L}(H)$. The latter is defined as weakest topology such that all linear maps

$$\begin{cases} \mathrm{L}(H) & \longrightarrow & \mathbf{C} \\ T & \mapsto & Tv \end{cases}$$

for $v \in H$ are continuous with respect to this topology. This means that a basis of neighborhoods of $T_0 \in \mathrm{L}(H)$ for this topology is given by a finite intersection

$$\bigcap_{1 \leqslant i \leqslant m} V_i$$

where

$$V_i = \{T \in \mathrm{L}(H) \mid \|Tv_i - T_0 v_i\| < \varepsilon\}$$

for some unit vectors $v_i \in H$ and some fixed $\varepsilon > 0$ (see, e.g., [**31**, p. 183]).

Example 3.3.5 (The regular representation on $L^2(\mathbf{R})$). Here is an example of infinite-dimensional unitary representation. We consider the additive group $\mathbf{R}$ of real numbers, with the usual topology, and the space $H = L^2(\mathbf{R})$, the space of square-integrable functions on $\mathbf{R}$, defined using Lebesgue measure. Defining

$$\varrho(t)f(x) = f(x + t),$$

we claim that we obtain a unitary representation. This is formally obvious, but as we have seen, the continuity requirement needs some care. We will see this in a greater context in Proposition 5.2.6 in Chapter 5, but we sketch the argument here.

First, one must check that the definition of $\varrho(t)f$ makes sense (since elements of $L^2(\mathbf{R})$ are not really functions): this is because if we change a measurable function f on a set of measure zero, we only change $x \mapsto f(x + t)$ on a translate of this set, which still have measure zero.

Then the unitarity of the action is clear: we have

$$\int_{\mathbf{R}} |\varrho(t)f(x)|^2 dx = \int_{\mathbf{R}} |f(x + t)|^2 dx = \int_{\mathbf{R}} |f(x)|^2 dx,$$

and only continuity remains to be checked. This is done in two steps using Proposition 3.3.3 (the reader can fill the outline, or look at Proposition 5.2.6). Fix $f \in L^2(\mathbf{R})$, and assume first that it is continuous and compactly supported. Then the continuity of $t \mapsto \varrho(t)f$ amounts to the limit

$$\lim_{h\to 0} \int_{\mathbf{R}} |f(x+t+h) - f(x+t)|^2 dx = 0$$

for all $t \in \mathbf{R}$, which follows from the dominated convergence theorem. Next, one uses the fact that continuous functions with compact support form a dense subspace of $L^2(\mathbf{R})$ (for the L^2-norm) in order to extend this continuity statement to all $f \in L^2(\mathbf{R})$.

The operations of representation theory, when applied to unitary representations, lead most of the time to other unitary representations. Here are the simplest cases:

Proposition 3.3.6 (Operations on unitary representations). *Let G be a topological group.*

(1) *If ϱ is a unitary representation of G, then any subrepresentation and any quotient representation are naturally unitary. Similarly, the restriction of ϱ to any subgroup H is a unitary representation with the same inner product.*

(2) *Direct sums of unitary representations are unitary with inner product*

$$\langle v_1 \oplus w_1, v_2 \oplus w_2\rangle_{\varrho_1\oplus\varrho_2} = \langle v_1, v_2\rangle_{\varrho_1} + \langle w_1, w_2\rangle_{\varrho_2}.$$

for $\varrho_1 \oplus \varrho_2$, so that the subrepresentations ϱ_1 and ϱ_2 in $\varrho_1 \oplus \varrho_2$ are orthogonal complements of each other.

(3) *The tensor product $\varrho_1 \otimes \varrho_2$ of finite-dimensional unitary representations ϱ_1 and ϱ_2 is unitary, with respect to the inner product defined for pure tensors by*

$$\langle v_1 \otimes w_1, v_2 \otimes w_2\rangle_{\varrho_1\otimes\varrho_2} = \langle v_1, v_2\rangle_{\varrho_1}\langle w_1, w_2\rangle_{\varrho_2}. \tag{3.4}$$

Similarly, external tensor products of finite-dimensional unitary representations, are unitary, with the same inner product on the underlying tensor-product space.

We leave to the reader the simple (and standard) argument that checks that the definition of (3.4) does extend to a well-defined inner product on the tensor product of two finite-dimensional Hilbert spaces.

Note that one can extend this to situations involving infinite-dimensional representations; this is very easy if either ϱ_1 or ϱ_2 is finite-dimensional, but it becomes rather tricky to define the tensor product of two infinite-dimensional Hilbert spaces. Similarly, extending the notion of induction to unitary representations of topological groups is not obvious; we will consider this in Chapter 5 for compact groups. In both cases, the reader may look at [**1**, App. A, App. E] for some more results and general facts.

Example 3.3.7. A special case of (3) concerns the tensor product of an arbitrary unitary representation $\varrho : G \longrightarrow \mathrm{U}(H)$ with a one-dimensional unitary representation $\chi : G \longrightarrow \mathbf{S}^1$. Indeed, in that case, one can define $\varrho \otimes \chi$ on the same Hilbert space H, by

$$(\varrho \otimes \chi)(g)v = \chi(g)\varrho(g)v,$$

with the same inner product, which is still invariant for $\varrho \otimes \chi$: for any v, $w \in H$, we have

$$\begin{aligned}\langle(\varrho \otimes \chi)(g)v, (\varrho \otimes \chi)(g)w\rangle_H &= \langle\chi(g)\varrho(g)v, \chi(g)\varrho(g)w\rangle_H \\ &= |\chi(g)|^2\langle\varrho(g)v, \varrho(g)w\rangle_H = \langle v, w\rangle_H.\end{aligned}$$

Now we discuss the contragredient in the context of unitary representations. There are special features here, which arise from the canonical duality properties of Hilbert spaces. For a Hilbert space H, the *conjugate space* $\bar{H}$ is defined to be the Hilbert space with the same underlying abelian group (i.e., addition) H, but with scalar multiplication and inner products defined by

$$\alpha \cdot v = \bar{\alpha} v, \qquad \langle v, w\rangle_{\bar{H}} = \langle w, v\rangle_H.$$

The point of the conjugate Hilbert space is that it is canonically isometric[4] (as Banach space) to the dual of H, by the map

$$\Phi \begin{cases} \bar{H} & \longrightarrow & H' \\ w & \mapsto & (\lambda_w \,:\, v \mapsto \langle v, w\rangle) \end{cases}$$

(vectors in w are "the same" as vectors in H, but because $\lambda_{\alpha w} = \bar{\alpha}\lambda_w$, this map is only linear when the conjugate Hilbert space structure is used as the source.)

If ϱ is a unitary representation of a group G on H, this allows us to rewrite the basic matrix coefficients $f_{v,\lambda}$ of ϱ using inner products on H only: given $\lambda = \lambda_w \in H'$ and $v \in H$, we have

$$f_{v,\lambda}(g) = \lambda_w(\varrho(g)v) = \langle \varrho(g)v, w\rangle.$$

These are now parametrized by the two vectors v and w in H; though it is formally better to see w as being an element of $\bar{H}$, one can usually dispense with this extra formalism without much loss.

Using the map Φ, we can also "transport" the contragredient representation of ϱ to an isomorphic representation $\bar{\varrho}$ acting on $\bar{H}$. Its character, when ϱ is finite-dimensional, is given by

$$\chi_{\bar{\varrho}}(g) = \overline{\chi_\varrho(g)}$$

(since the eigenvalues of a unitary matrix are roots of unity, hence their inverse is the same as their conjugate; see Lemma 3.3.10 below also).

Exercise 3.3.8 (Matrix coefficients of the conjugate representation). Let ϱ be a unitary representation of G on a Hilbert space H. Show that the functions

$$g \mapsto \overline{\langle \varrho(g)v, w\rangle},$$

for v, $w \in H$ are matrix coefficients of $\bar{\varrho}$.

Example 3.3.9 (A unitarizability criterion). Let $G = \mathbf{R}$ with its usual topology. We have seen that there are many one-dimensional representations of G as a topological group, given by

$$\omega_s \,:\, \begin{cases} \mathbf{R} & \longrightarrow & \mathbf{C}^\times \\ x & \mapsto & e^{sx} \end{cases}$$

for $s \in \mathbf{C}$ (these are different functions, hence non-isomorphic representations).

However, these are only unitary (or, more properly speaking, unitarizable) when $s = it$ is purely imaginary. Indeed, when $\mathrm{Re}(s) = 0$, we have $|\omega_s(x)| = 1$ for all $x \in \mathbf{R}$, and the unit circle is the unitary group of the 1-dimensional Hilbert space $\mathbf{C}$ with inner product $z\bar{w}$. Conversely, the following lemma is a quick convenient necessary condition for unitarity or unitarizability, because it gives a property which does not depend on any information on the inner product, and it implies that ω_s is not unitarizable otherwise.

Lemma 3.3.10. *Let ϱ be a finite-dimensional unitary representation of a group G. Then any eigenvalue of $\varrho(g)$, for any g, is a complex number of modulus* 1.

[4] This is the Riesz representation theorem for Hilbert spaces.

Proof. This is linear algebra for unitary matrices, of which the $\varrho(g)$ are examples... □

This lemma applies also to the representations ϱ_m of $\mathrm{SL}_2(\mathbf{C})$ of Examples 2.6.1 and 2.7.38: from (2.37) – for instance – we see that if $m \geqslant 1$, ϱ_m is not unitarizable (since the restriction to the subgroup T is diagonal with explicit eigenvalues which are not of modulus 1). However, the restriction of ϱ_m to the compact subgroup $\mathrm{SU}_2(\mathbf{C})$ is unitarizable.

Along the same lines, observe that if we consider the compact group $\mathbf{R}/\mathbf{Z}$, its one-dimensional representations induce, by composition, some representations of $\mathbf{R}$

$$\mathbf{R} \longrightarrow \mathbf{R}/\mathbf{Z} \longrightarrow \mathbf{C}^\times,$$

which must be among the ω_s. Which ones occur in this manner is easy to see: we must have $\mathbf{Z} \subset \ker(\omega_s)$, and from $\omega_s(1) = 1$, it follows that $s = 2ik\pi$ for some *integer* $k \in \mathbf{Z}$. In particular, we observe a feature which will turn out to be quite general: all these representations of $\mathbf{R}/\mathbf{Z}$ are unitary!

Example 3.3.11 (Regular representation of a finite group). Let G be a finite group, and $C(G)$ the space of complex-valued functions on G, with the regular representation of G. A natural inner product on the vector space $C(G)$ is

$$\langle \varphi_1, \varphi_2 \rangle = \frac{1}{|G|} \sum_{x \in G} \varphi_1(x)\overline{\varphi_2(x)}$$

(one could omit the normalizing factor $1/|G|$, but it has the advantage that $\|1\| = 1$ for the constant function[5] 1 on G, independently of the order of G.)

It is quite natural that, with respect to this inner product, the representation $C(G)$ is unitary. Indeed, we have

$$\| \operatorname{reg}(g)\varphi\|^2 = \sum_{x \in G} |\varphi(xg)|^2 = \sum_{y \in G} |\varphi(y)|^2 = \|\varphi\|^2$$

for all $g \in G$, using the change of variable $y = xg$.

A similar property holds for all locally compact groups, but if G is infinite, the inner product must be defined using integral on G with respect to a natural measure μ, and the space $C(G)$ must be replaced by the Hilbert space $L^2(G, \mu)$. We will come back to this later (see Proposition 5.2.6).

In addition to the usual formalism of direct sum, the extra structure of Hilbert spaces leads to a definition of infinite orthogonal direct sums. If G is a topological group and $(\varrho_i)_{i \in I}$ is any family of unitary representations of G, acting on the Hilbert spaces H_i, we can define the Hilbert space orthogonal direct sum

$$H = \bigoplus_{i \in I} H_i, \tag{3.5}$$

and a corresponding representation $\varrho = \oplus_i \varrho_i$ acting on H_i. Precisely, recall that H is defined to be the space of families $v = (v_i)_{i \in I}$ such that

$$\|v\|_H^2 = \sum_{i \in I} \|v_i\|_i^2 < +\infty,$$

and we define

$$\varrho(g)v = (\varrho_i(g)v_i)_{i \in I},$$

[5] This should not be confused with the neutral element in the group.

which is easily checked to be a unitary representation acting on H (see Exercise 3.3.12 below). Of course, for each i, the subspace $H_i \subset H$ is a subrepresentation of H isomorphic to ϱ_i. Moreover, the "standard" direct sum of the H_i (the space of families (v_i) where v_i is zero for all but finitely many i) is a dense subspace of H. It is stable under the action of ϱ, but since it is not closed in general, it is usually not a subrepresentation in the topological sense.

Exercise 3.3.12 (Pre-unitary representation). It is often convenient to define a unitary representation by first considering an action of G on a pre-Hilbert space, which "extends by continuity" to a proper unitay representation (e.g., when defining a representation of a space of functions, it may be easier to work with a dense subspace of regular functions; for a concrete example, see the construction of the regular representation of a compact topological group in Proposition 5.2.6)). We consider a fixed topological group G.

A pre-unitary representation of G is a strongly continuous homomorphism

$$\varrho : G \longrightarrow \mathrm{U}(H_0),$$

where H_0 is a pre-Hilbert space, i.e., a complex vector space given with a (positive-definite) inner product, but which is not necessarily complete.

(1) Show that if ϱ is a pre-unitary representation, the operators $\varrho(g)$ extend by continuity to unitary operators of the completion H of H_0, and that the resulting map is a unitary representation of G, such that $H_0 \subset H$ is a stable subspace. [Hint: To check the strong continuity, use the fact that H_0 is dense in H.]

(2) Suppose H is a Hilbert space and $H_0 \subset H$ a dense subspace. If ϱ is a pre-unitary representation on H_0, explain why the resulting unitary representation is a representation of G on H.

(3) Use this to check that the Hilbert direct sum construction of (3.5) above leads to unitary representations.

Unitary representations are better behaved than general (complex) representations. One of the main reasons is the following fact:

Proposition 3.3.13 (Reducibility of unitary representations). *Let $G \xrightarrow{\varrho} \mathrm{U}(H)$ be a unitary representation of a topological group G. Then any closed subspace $F \subset H$ invariant under ϱ has a stable closed complement given by $F^{\perp} \subset H$. In particular, any finite-dimensional unitary representation is semisimple.*

Proof. Since any finite-dimensional subspace of a Hilbert space is closed, the second part follows from the first using the criterion of Lemma 2.2.8.

Thus we consider a subrepresentation $F \subset H$. Using the inner product, we can very easily construct a stable complement: we may just consider the orthogonal complement

$$F^{\perp} = \{v \in H \mid \langle v, w\rangle = 0 \text{ for all } w \in F\} \subset H.$$

Indeed, the theory of Hilbert spaces[6] shows that $F^{\perp}$ is closed in H and that $F \oplus F^{\perp} = H$. From the fact that ϱ preserves the inner product, the same property follows for its orthogonal complement: if $v \in F^{\perp}$, we have

$$\langle \varrho(g)v, w\rangle = \langle v, \varrho^{-1}(g)w\rangle = 0$$

for all $g \in G$ and $w \in F$, i.e., $F^{\perp}$ is indeed a subrepresentation of H. □

[6] If H is finite-dimensional, of course, this is mere linear algebra.

Example 3.3.14 (Failure of semisimplicity for unitary representations). Although this property is related to semisimplicity, it is *not* the case that any unitary representation

$$\varrho : G \longrightarrow \mathrm{U}(H)$$

is semisimple, even in the sense that there exists a family $(H_i)_{i\in I}$ of stable subspaces of H such that

$$H = \bigoplus_{i\in I} H_i$$

(in the Hilbert-space sense described above). The reader can of course look back at the proof of Lemma 2.2.8 where the equivalence of semisimplicity and complete reducibility was proved for the "algebraic" case: the problem is that the first step, the *existence* of an irreducible subrepresentation of H (which is Exercise 2.2.10), may fail in this context. To see this, take the representation ϱ of $\mathbf{R}$ on $L^2(\mathbf{R})$ described in Example 3.3.5. This is of course infinite-dimensional *but $L^2(\mathbf{R})$ contains no irreducible subrepresentation*! We can not quite prove this rigorously yet, but the following explains what happens: first, because $\mathbf{R}$ is abelian, all its irreducible unitary representations are of dimension 1 (as is the case for finite abelian groups, though the proof is harder here since one must exclude possible infinite-dimensional representations), and then, by Proposition 3.2.3 and the unitarizability criterion, these are given by

$$\chi_x \begin{cases} \mathbf{R} & \longrightarrow & \mathbf{C}^\times \\ t & \mapsto & e^{itx} \end{cases}$$

for $x \in \mathbf{R}$. Now, a non-zero function $f \in L^2(\mathbf{R})$ spans an irreducible subrepresentation of ϱ isomorphic to χ_x if and only if we have

$$f(x+t) = \varrho(t)f(x) = \chi_x(t)f(x) = e^{itx}f(x)$$

for all x and $t \in \mathbf{R}$. But this means that $|f(t)| = |f(0)|$ is constant for all $t \in \mathbf{R}$, and this constant is non-zero since we started with $f \neq 0$. However, we get

$$\int_{\mathbf{R}} |f(x)|dx = |f(0)| \int_{\mathbf{R}} dx = +\infty,$$

which contradicts the assumption $f \in L^2(\mathbf{R})$...

In Chapter 5, we will see that this type of behavior does not occur for *compact* topological groups. But this shows, obviously, that the study of unitary representations of non-compact groups will be frought with new difficulties...

Along the same lines, the following related result is also very useful:

Lemma 3.3.15 (Unrelated unitary subrepresentations are orthogonal). *Let G be a topological group and let $G \xrightarrow{\varrho} \mathrm{U}(H)$ be a unitary representation. If H_1 and H_2 are subrepresentations of H such that there is no non-zero G-intertwiner $H_1 \rightarrow H_2$, then H_1 and H_2 are orthogonal. In particular, isotypic components in H of non-isomorphic irreducible representations of G are pairwise orthogonal.*

Proof. Consider the orthogonal projector

$$\Phi : H \longrightarrow H$$

on H_2. This linear map Φ is also a G-homomorphism because of the previous proposition: if $v \in H$ and $g \in G$, its projection $\Phi(v)$ is characterized by

$$v = \Phi(v) + (v - \Phi(v)), \qquad \Phi(v) \in H_1, \quad v - \Phi(v) \in H_1^{\perp},$$

and the subsequent relation

$$\varrho(g)v = \varrho(g)(\Phi(v)) + \varrho(g)(v - \Phi(v))$$

together with the condition $\varrho(g)\Phi(v) \in H_1$, $\varrho(g)(v - \Phi(v)) \in H_1^{\perp}$ (which follow from the fact that H_1 is a subrepresentation) imply that

$$\Phi(\varrho(g)v) = \varrho(g)\Phi(v),$$

which gives the desired property.

Since the image of Φ is H_2, its restriction to H_1 is a linear map

$$H_1 \longrightarrow H_2$$

which is a G-intertwiner. The assumption then says that it is zero, so that $H_1 \subset \ker(\Phi) = H_2^{\perp}$, which is the same as to say that H_1 and H_2 are orthogonal.

The last statement is of course a corollary of this fact together with Schur's Lemma. □

CHAPTER 4

Linear representations of finite groups

In this chapter, we take up the special case of finite groups, building on the basic results of Section 2.7. There are however still two very distinct cases: if the field k has characteristic coprime with the order of a finite group G (for instance, if $k = \mathbf{C}$, or any other field of characteristic 0), a fundamental result of Maschke shows that *any k-representation of G is semisimple*. Thus, we can use characters to characterize all (finite-dimensional) representations of G, and this leads to very powerful methods to analyze representations. Most of this chapter will be devoted to this case. However, if k is of characteristic p dividing the order of G, the semisimplicity property fails. The classification and structure of representations of G is then much more subtle; since the author knows next to nothing about this case, we will only give some examples and general remarks in Section 4.8.2.

4.1. Maschke's Theorem

As already hinted, the next result is the most important result about the representation theory of finite groups:

THEOREM 4.1.1 (Maschke). *Let G be a finite group, and let k be a field with characteristic* not dividing $|G|$. *Then any k-linear representation*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of G is semisimple. In fact, the converse is also true: if all k-representations of G are semisimple, then the characteristic of k does not divide $|G|$.

Thus, in some sense, in the case where Maschke's Theorem applies, the classification of all representations of G is reduced to the question of classifying the irreducible ones. Note that it is not required to assume that k be algebraically closed here.

PROOF. We use the criterion of Lemma 2.2.8. For a given subrepresentation $F \subset E$, the idea is to construct a stable supplement $F^{\perp}$ as the kernel of a linear projection

$$P : E \longrightarrow E$$

with image F which is a G-morphism, i.e., $P \in \mathrm{Hom}_G(E, E)$. Indeed, if $P^2 = P$ (which means P is a projection) and $\mathrm{Im}(P) = F$, we have

$$E = F \oplus \ker(P),$$

and of course $\ker(P)$ is a subrepresentation if $P \in \mathrm{Hom}_G(E, E)$.

From linear algebra again, we know the existence of a projection $p \in \mathrm{Hom}_k(E, E)$ with $\mathrm{Im}(p) = F$, but a priori not one that commutes with the action of G. Note that $p \in \mathrm{Hom}_G(E, E)$ means that $p \in \mathrm{End}_k(E)^G$ (see (2.18)). The trick is to construct P using p by *averaging* the action (2.16) of G on p in order to make it invariant.

Let then

$$P = \frac{1}{|G|} \sum_{g \in G} g \cdot p \in \mathrm{End}_k(E).$$

We claim that P is the desired projection in $\mathrm{End}_G(E)$. The first thing to notice is that it is this definition which requires that $p \nmid |G|$, since $|G|$ must be invertible in k in order to compute P.

By averaging, it is certainly the case that $P \in \mathrm{End}_G(E) = \mathrm{End}_k(E)^G$: acting on the left by some $h \in G$ just permutes the summands

$$h \cdot P = \frac{1}{|G|} \sum_{g \in G} h \cdot (g \cdot p) = \frac{1}{|G|} \sum_{g \in G} (hg) \cdot p = \frac{1}{|G|} \sum_{x \in G} x \cdot p = P$$

(in the notation of Example 3.1.1, we are using the fact that $P = e \cdot p$ where $e = \frac{1}{|G|} s \in k(G)$, and that $he = e$ by (3.2)).

Next, $\mathrm{Im}(P) \subset \mathrm{Im}(p) = F$: indeed, F is G-invariant and each term

$$(g \cdot p)v = \varrho(g)(p(\varrho(g^{-1})v)) \in F$$

in the sum is in F for any fixed v. Moreover, P is the identity on F, since p is and F is stable: for $v \in F$, we have

$$P(v) = \frac{1}{|G|} \sum_{g \in G} \varrho(g) p(\varrho(g^{-1})v) = \frac{1}{|G|} \sum_{g \in G} \varrho(gg^{-1})v = v.$$

Thus $(P \circ P)(v) = P(P(v)) = P(v)$ since $P(v) \in F$, and hence P is indeed an intertwining projection onto F.

We now prove the converse of Maschke's Theorem. In fact, the result is stronger than what we claim: we now show that $C_k(G)$ is never semisimple if k has characteristic p dividing $|G|$. To do this, we consider the subspace

$$C_0 = \{\varphi \in C_k(G) \mid \sum_{g \in G} \varphi(g) = 0\} \subset C_k(G).$$

This subspace is always a subrepresentation of $C_k(G)$, as one checks immediately (as before, the values of $\mathrm{reg}(g)\varphi$ are a permutation of the values of φ). As the kernel of the non-zero linear form

$$\lambda : \varphi \mapsto \sum_{g \in G} \varphi(g),$$

we see that C_0 is in fact of codimension 1 in $C_k(G)$. If $p \nmid |G|$, a complementary stable subspace is easy to find: it is the space of constant functions (it is in fact the unique stable complement). But if $\varphi = c \in k$ is constant, and $p \mid |G|$, we have $\lambda(\varphi) = c|G| = 0$, so this complement does not work in characteristic p. We now check that no other will do: if φ_0 is a basis of such a complement, the action of G on $k\varphi_0$ is by a one-dimensional representation ϱ, so we have

$$\mathrm{reg}(g)\varphi_0 = \varrho(g)\varphi_0$$

for all $g \in G$; evaluating at 1, we find that $\varphi_0(g) = \varrho(g)\varphi_0(1)$ for all G. But now

$$\lambda(\varphi_0) = \varphi_0(1) \sum_{g \in G} \varrho(g).$$

The last sum, which is $\lambda(\varrho)$ is, however, equal to 0, and this is a contradiction. Indeed, either ϱ is trivial, and then the value is $|G| = 0$ in k, or there exists $x \in G$ with $\varrho(x) \neq 1$, and then writing

$$\lambda(\varrho) = \sum_{g \in G} \varrho(g) = \sum_{h \in G} \varrho(xh) = \varrho(x)\lambda(\varrho)$$

implies that $\lambda(\varrho) = 0$ also! □

Remark 4.1.2 (Semisimplicity of unitary representations). If $k = \mathbf{C}$, we can also prove Theorem 4.1.1 by exploiting Proposition 3.3.13, at least when dealing with finite-dimensional representations. Indeed, we have the following result:

Proposition 4.1.3 (Unitarizability for finite groups). *For a finite group G, any finite-dimensional representation of G over $\mathbf{C}$ is unitarizable.*

Proof. The idea is similar to the one in the proof of Maschke's Theorem. Indeed, let ϱ be a finite-dimensional representation of G on E. What must be done is to find an inner product $\langle \cdot, \cdot \rangle$ on E with respect to which ϱ is unitary, i.e., such that

$$\langle \varrho(g)v, \varrho(g)w \rangle = \langle v, w \rangle$$

for all v, $w \in E$. Now, since E is finite-dimensional, we can certainly find some inner product $\langle \cdot, \cdot \rangle_0$ on E, although it is not necessarily invariant. But then if we let

$$\langle v, w \rangle = \frac{1}{|G|} \sum_{g \in G} \langle \varrho(g)v, \varrho(g)w \rangle_0,$$

it is easy to check that we obtain the desired invariant inner product. □

More generally, suppose $\langle \cdot, \cdot \rangle_0$ is a non-negative, but not necessarily positive-definite, hermitian form on E, with kernel

$$F = \{v \in E \mid \langle v, v \rangle = 0\}.$$

Then it is clear that the construction of $\langle \cdot, \cdot \rangle$ above still leads to an invariant non-negative hermitian form. By positivity, it will be a genuine, positive-definite, inner product if

$$\bigcap_{g \in G} g \cdot F = 0.$$

An example of this is the regular representation, where we can take

$$\langle \varphi_1, \varphi_2 \rangle_0 = \varphi_1(1)\overline{\varphi_2(1)}.$$

This has a huge kernel F (all functions vanishing at 1) but since

$$\mathrm{reg}(g)F = \{\varphi \in C(G) \mid \mathrm{reg}(g)\varphi(1) = \varphi(g) = 0\},$$

the intersection of the translates of F is in fact 0. Rather naturally, the resulting inner product on $C(G)$ is the same described in Example 3.3.11.

The meaning of Maschke's Theorem is that for any k-representation ϱ of G on a vector space E, if $|G|$ is invertible in k, there is a direct sum decomposition of E in irreducible stable subspaces. As already discussed in Chapter 2, this decomposition is not unique. However, by Proposition 2.7.7, the isotypic components of ϱ, denoted $M(\pi)$ or $M_E(\pi)$, are defined for any irreducible k-representation π of G, and they are intrinsic subspaces of E such that

$$E = \bigoplus_{\pi} M_E(\pi), \tag{4.1}$$

where π runs over isomorphism classes of irreducible k-representations of G. Recall that because they are intrinsic, it follows that for any G-homomorphism

$$\Phi : E \longrightarrow F,$$

the restriction of Φ to $M_E(\pi)$ gives a linear map

$$M_E(\pi) \longrightarrow M_F(\pi).$$

In order to analyze the representations concretely, one needs some way of obtaining information concerning this decomposition; we will see how to describe (when k is algebraically closed) explicitly the projectors on E mapping onto the isotypic components, and when k is of characteristic 0, how to use characters to compute the multiplicities of the irreducible representations (which are of course related to the dimension of the isotypic components.)

4.2. Applications of Maschke's Theorem

We are now going to apply Maschke's Theorem. First of all, here and in the rest of this chapter (up to Section 4.8.2), unless otherwise indicated, we *assume that k is algebraically closed and that the characteristic of k does not divide $|G|$,* so that Maschke's Theorem is applicable, as well as Schur's Lemma 2.7.13.

Applying first Maschke's Theorem to the regular representation of G on $C_k(G)$, we deduce from Corollary 2.7.26 a fundamental result:

Corollary 4.2.1 (Decomposition of the regular representation). *Let G be a finite group and let k be an algebraically closed field of characteristic not dividing $|G|$. Then the regular representation of G on $C_k(G)$ is isomorphic to the direct sum, over all irreducible representations ϱ of G, up to isomorphism, of subrepresentations isomorphic to* $\dim(\varrho)$ *copies of ϱ.*

In particular, we have

$$\sum_{\varrho} (\dim \varrho)^2 = |G|, \tag{4.2}$$

where the sum is over the set of isomorphism classes of representations of G, which can be identified with the set of characters of irreducible representations.

One naturally wants to get more information about the irreducible representations than what is contained in the formula (4.2). The first basic question is: what is the *number* of irreducible representations (up to isomorphism)? The general answer is known, but we start with a particularly simple case:

Proposition 4.2.2 (Irreducible representations of finite abelian groups). *Let G be a finite group and k an algebraically closed field of characteristic not dividing $|G|$. Then all irreducible finite-dimensional representations of G are of dimension* 1 *if and only if G is commutative. In particular, there are $|G|$ non-isomorphic irreducible k-representations of G.*

Proof of Proposition 4.2.2. We know that the one-dimensional representations of G are in bijection with those of the abelianized group $G/[G,G]$ (Proposition 2.6.6). Thus if all irreducible k-representations of G are of dimension 1, Corollary 4.2.1 implies that $|G| = |G/[G,G]|$ (the left-hand sides being equal for G and $G/[G,G]$), which means that $[G,G] = 1$, i.e., that G is commutative. □

Although the representations of abelian groups are quite elementary in comparison with the general case, they are of great importance in applications. We will say more about them in Section 4.5, which includes in particular a sketch of the proof of Dirichlet's Theorem on primes in arithmetic progressions. (That later section could be read right now without much difficulty.)

Note that one can not remove the assumption on k: there are cases where G is non-abelian, k is algebraically closed of characteristic dividing G, and the only irreducible

k-representation of G is trivial (indeed, this is true for all non-abelian groups of prime order p and algebraically closed fields of characteristic p, see, e.g., [**8**, 27.28]; examples are given by the groups from Example 2.7.33.)

EXAMPLE 4.2.3. Consider $G = \mathfrak{S}_3$, the symmetric group on 3 letters. It is non-abelian of order 6, and hence the only possible values for the degrees of irreducible $\mathbf{C}$-representations of G are 1, 1 and 2 (there are no other integers with squares summing to 6, where not all are equal to 1). The two one-dimensional representations are of course the trivial one and the signature

$$\varepsilon : \mathfrak{S}_3 \longrightarrow \{\pm 1\} \subset \mathbf{C}^\times,$$

and the 2-dimensional one is isomorphic to the representation by permutation of the coordinates on the vector space

$$E = \{(x, y, z) \in \mathbf{C}^3 \mid x + y + z = 0\}.$$

More generally, the decomposition of the regular representation implies that there are at most $|G|$ irreducible representations, and Proposition 4.2.2 shows that this upper bound is reached if and only G is abelian. In fact, we have the following:

THEOREM 4.2.4 (Number of irreducible characters). *Let G be a finite group, k an algebraically closed field of characteristic not dividing $|G|$. Then the number of irreducible k-representations of G is equal to the number of conjugacy classes in G.*

This is another striking fact. This tells us immediately, for instance, that the symmetric group $\mathfrak{S}_{24}$ has exactly 1575 irreducible complex representations up to isomorphism. Indeed, the number of conjugacy classes in $\mathfrak{S}_n$ is, via the cycle type of permutations, the same as the number $p(n)$ of partitions of n, i.e., the number of solutions $(r_1, \ldots, r_n)$ of the equation

$$n = 1 \cdot r_1 + 2 \cdot r_2 + \cdots + n \cdot r_n,$$

in non-negative integers (in this bijection, $(r_1, \ldots, r_n)$ corresponds to the permutations which have cycle decomposition with r_1 fixed points, r_2 disjoint transpositions, ..., and r_n cycles of length n.) Of course, it might not be obvious that $p(24) = 1575$, but computers are here to confirm this. We will give more comments on this theorem after its proof.

PROOF. We have the decomposition

$$C_k(G) = \bigoplus_{\varrho} M(\varrho), \tag{4.3}$$

where $M(\varrho)$ is the space spanned by all matrix-coefficients of the irreducible k-representation ϱ. To compute the number of summands, we try to find some invariant of $C_k(G)$ which will involve "counting" each ϱ only once. The dimension does not work since $\dim M(\varrho) = \dim(\varrho)^2$; computing the intertwiners from $C_k(G)$ to itself is also tempting but since $M(\varrho)$ is a direct sum of $\dim(\varrho)$ copies of ϱ, we have (by Schur's Lemma) again

$$\dim \mathrm{Hom}_G(C_k(G), C_k(G)) = \sum_{\varrho} \dim \mathrm{Hom}_G(M(\varrho), M(\varrho)) = \sum_{\varrho} (\dim \varrho)^2 = |G|.$$

The way (or one way at least) to do this turns out to be a bit tricky: we use the fact that $C_k(G)$ carries in fact a representation π of $G \times G$ defined by

$$\pi(g_1, g_2)f(x) = f(g_1^{-1} x g_2),$$

and that the decomposition (4.3) is in fact a decomposition of $C_k(G)$ into subrepresentations of π; indeed, if $f_{v,\lambda} \in M(\varrho)$ is a matrix coefficient

$$f_{v,\lambda}(x) = \langle \lambda, \varrho(g)v \rangle$$

of an irreducible representation ϱ, we have

$$\pi(g_1, g_2)f_{v,\lambda}(x) = \langle \lambda, \varrho(g_1^{-1}xg_2)v \rangle = \langle \tilde{\varrho}(g_1)\lambda, \varrho(x)\varrho(g_2)v \rangle = f_{\tilde{\varrho}(g_1)\lambda, \varrho(g_2)v}(x)$$

so that $M(\varrho)$ is indeed stable under the action of $G \times G$.

Now the point is that $M(\varrho)$, as a subrepresentation of π in $C_k(G)$, is isomorphic to the external tensor product $\tilde{\varrho} \boxtimes \varrho$. Indeed, if ϱ acts on the space E, this isomorphism is the canonical one given by

$$\begin{cases} E' \otimes E & \longrightarrow & M(\varrho) \\ \lambda \otimes v & \mapsto & f_{v,\lambda} \end{cases}$$

which is a linear isomorphism by linear independence of the matrix coefficients, and an intertwiner by the computation just performed.[1]

Since the representations $\tilde{\varrho} \boxtimes \varrho$ of $G \times G$ are all irreducible and non isomorphic, as ϱ runs over the irreducible representations of G, by Proposition 2.3.17, each appears with multiplicity 1 in π. As a consequence, we can use Schur's Lemma (on $G \times G$) to express the number of ϱ by the formula

$$\dim \mathrm{Hom}_{G\times G}(\pi, \pi) = \sum_{\varrho_1, \varrho_2} \dim \mathrm{Hom}_{G\times G}(\varrho_1 \boxtimes \tilde{\varrho}_1, \varrho_2 \boxtimes \tilde{\varrho}_2) = \sum_{\varrho} 1.$$

We now compute directly the left-hand side dimension to deduce the desired formula. In fact, consider a linear map

$$\Phi \,:\, C_k(G) \longrightarrow C_k(G)$$

which commutes with the representation π. If $\delta_g \in C_k(G)$ denotes the function which is 0 except at $x = g$, where it is equal to 1, and

$$\ell_g = \Phi(\delta_g),$$

we must therefore have

$$\pi(g_1, g_2)\ell_g = \Phi(\pi(g_1, g_2)\delta_g) = \Phi(\delta_{g_1 g g_2^{-1}}) = \ell_{g_1 g g_2^{-1}}$$

for all g, g_1, $g_2 \in G$. Evaluating at $x \in G$ means that we must have

$$\ell_g(g_1^{-1}xg_2) = \ell_{g_1 g g_2^{-1}}(x), \tag{4.4}$$

and in fact, since (δ_g) is a basis of $C_k(G)$, these equations on ℓ_g are equivalent with Φ being an intertwiner of π with itself.

We can solve these equations as follows: picking first $g_2 = 1$ and $g_1 = x^{-1}$, and then $g_1 = 1$ and $g_2 = x$, we get quickly the two relations

$$\ell_{gx^{-1}}(1) = \ell_g(x) = \ell_{x^{-1}g}(1). \tag{4.5}$$

Thus Φ is entirely defined by the function $\psi \,:\, G \longrightarrow k$ defined by

$$\psi(g) = \ell_g(1).$$

In view of the two expressions for $\ell_g(x)$ above, this function must satisfy

$$\psi(ab) = \psi(ba) \tag{4.6}$$

[1] Note that the order of the factors is important here! If we consider $E \otimes E'$ instead, we do not obtain an intertwiner...

for all a, $b \in G$. But conversely, is a function ψ satisfies this relation, defining $\ell_g(x)$ by

$$\ell_g(x) = \psi(gx^{-1}) = \psi(x^{-1}g),$$

(see (4.5)), we obtain

$$\ell_g(g_1^{-1}xg_2) = \psi(gg_2^{-1}x^{-1}g_1)$$

and

$$\ell_{g_1gg_2^{-1}}(x) = \psi(x^{-1}g_1gg_2^{-1}),$$

from which (4.4) follows by putting $a = gg_2^{-1}$, $b = x^{-1}g_1$ in (4.6).

The conclusion is that $\mathrm{Hom}_{G\times G}(\pi, \pi)$ is isomorphic (by mapping Φ to ψ) to the linear space $c_k(G)$ of all functions ψ such that (4.6) holds. But this condition is equivalent with

$$\psi(x) = \psi(gxg^{-1}), \qquad \text{for all } x, g \in G,$$

or in other words, $c_k(G)$ is the space of class-functions on G. Since the dimension of $c_k(G)$ is equal to the number of conjugacy classes of G, by definition (a class function is determined by the values at the conjugacy classes), we obtain the desired formula. □

Remark 4.2.5 (Sum of dimensions). Thus we have "directly accessible" group-theoretic expressions for the number of irreducible k-representations of a finite group G (which is the number of conjugacy classes), and for the sum of the squares of their degrees (which is simply $|G|$). It seems natural to ask: what about the sum of the degrees themselves, or what about other powers of $\dim(\varrho)$? Although there does not seem to exist any nice expression valid for all groups, there are some special cases (including important examples like symmetric groups) where the sum of the dimensions has a nice interpretation, as explained in Lemma 6.2.6 (in Chapter 6).

Remark 4.2.6 (Bijections, anyone?). Theorem 4.2.4 is extremely striking; since it gives an equality between the cardinality of the set of irreducible characters of G (over an algebraically closed field of characteristic not dividing $|G|$) and the cardinality of the set of conjugacy classes, it implies that there exist some bijection between the two sets. However, even for $k = \mathbf{C}$, there is no general natural definition of such a bijection, and there probably is none.

Despite this, there are many cases where one understand both the conjugacy classes and the irreducible representations, and where some rough features of the two sets seem to correspond in tantalizing parallels (see the discussion of $\mathrm{GL}_2(\mathbf{F}_p)$ later). And in some particularly favorable circumstances, a precise correspondence can be found; the most striking and important of these cases is that of the symmetric groups $\mathfrak{S}_m$.

Another equally striking aspect of the result is that the number of irreducible representations does not depend on the field k. Here again, this means that there are bijections between the sets of irreducible characters for different fields. This is of course surprising when the characteristics of the fields are distinct! One may also ask here if such a bijection can be described explicitly, and the situation is better than the previous one. Indeed, although a completely canonical correspondence does not seem to be obtainable, Brauer developed a theory which – among other things – does lead to bijections between irreducible characters of G over any algebraically closed fields of characteristic coprime with $|G|$. See, e.g, [**19**, Ch. 15, Th. 15.13] or [**34**, Ch. 18], for an account of this theory, from which it follows, in particular, that the family of dimensions of the irreducible characters over two such fields always coincide.

4.3. Decomposition of representations

4.3.1. The projection on invariant vectors. Especially when written in terms of the group algebra (using Example 3.1.1), the core argument of the proof of Maschke's Theorem can be immediately generalized:

PROPOSITION 4.3.1 (Projection on the invariant vectors). *Let G be a finite group and k an algebraically closed field of characteristic not dividing $|G|$. For any k-representation $\varrho : G \longrightarrow \mathrm{GL}(E)$, the map*

$$\frac{1}{|G|}\sum_{g\in G}\varrho(g) \,:\, E \longrightarrow E$$

is a homomorphism of representations, which is a projection with image equal to E^G, the space of invariant vectors in E. Moreover, if E is finite-dimensional, we have

$$\dim(E^G) = \frac{1}{|G|}\sum_{g\in G}\chi_\varrho(g) \tag{4.7}$$

as equality in the field k.

PROOF. Let P be the indicated linear map. As in the proof of Maschke's Theorem, we see that that $\mathrm{Im}(P) \subset E^G$, and that P is the identity on E^G. Moreover, we have

$$P(\varrho(g)v) = P(v) = \varrho(g)P(v),$$

since $\mathrm{Im}(P) \subset E^G$. All this shows that $P \in \mathrm{Hom}_G(E,E)$ is a projection with image exactly equal to E^G.

Finally, the trace of a projection is equal to the dimension of the image, as seen in the field k, hence the last formula. □

We can use fruitfully Proposition 4.3.1 in many ways. One is to take any representation for which we know the invariants, and to see what form the projection takes. The next few sections give some important examples.

4.3.2. "Orthogonality" of characters and matrix coefficients. Consider G and k as before. If

$$\pi_1 \,:\, G \longrightarrow \mathrm{GL}(E_1), \qquad \pi_2 \,:\, G \longrightarrow \mathrm{GL}(E_2)$$

are irreducible k-representations, we know by Schur's Lemma that for the natural action of G on $E = \mathrm{Hom}_k(E_1, E_2)$, the space of invariants

$$\mathrm{Hom}_G(E_1, E_2) = \mathrm{Hom}_k(E_1, E_2)^G$$

has dimension 0 or 1 depending on whether π_1 and π_2 are isomorphic. Applying Proposition 4.3.1, this leads to some fundamental facts.

We assume, to begin with, that π_1 is *not* isomorphic to π_2. Then the invariant space is 0, and hence the associated projector is also zero: for any $\Phi : E_1 \longrightarrow E_2$, we have

$$\frac{1}{|G|}\sum_{g\in G} g\cdot\Phi = 0. \tag{4.8}$$

To see what this means concretely, we select some $\Phi \in E = \mathrm{Hom}_k(E_1, E_2)$. Because there is no intrinsic relation between the spaces here, some choice must be made. We consider linear maps of rank 1: let $\lambda \in E_1'$ be a linear form and $w \in E_2$ be a vector, and let

$$\Phi \,:\, v \mapsto \langle\lambda, v\rangle w$$

be the corresponding rank 1 map in $\mathrm{Hom}_k(E_1, E_2)$. Spelling out the identity (4.8) by applying it to a vector $v \in E_1$, we get

$$0 = \frac{1}{|G|} \sum_{g \in G} \pi_2(g)(\langle \lambda, \pi_1(g^{-1})v) \rangle w) = \frac{1}{|G|} \sum_{g \in G} f_{v,\lambda}(g) \pi_2(g^{-1})w$$

for all v (we replaced g by g^{-1} in the sum, which merely permutes the terms).

We can make this even more concrete by applying an arbitrary linear form $\mu \in E_2'$ to obtain numerical identities:

Corollary 4.3.2. *With notation as above, let π_1, π_2 be* non-isomorphic *irreducible representations of G; then for all vectors $v \in E_1$, $w \in E_2$ and linear forms $\lambda \in E_1'$, $\mu \in E_2'$, we have*

$$\frac{1}{|G|} \sum_{g \in G} f_{v,\lambda}(g) f_{w,\mu}(g^{-1}) = \frac{1}{|G|} \sum_{g \in G} \langle \lambda, g \cdot v \rangle_{E_1} \langle \mu, g^{-1} \cdot w \rangle_{E_2} = 0. \tag{4.9}$$

Because such sums with come out often, it is convenient to make the following definition:

Definition 4.3.3 ("Inner-product" of functions on G). Let G be a finite group and let k be a field with $|G|$ invertible in k. For φ_1, φ_2 in $C_k(G)$, we denote

$$[\varphi_1, \varphi_2] = \frac{1}{|G|} \sum_{g \in G} \varphi_1(g) \varphi_2(g^{-1}).$$

This is a non-degenerate[2] symmetric bilinear form on $C_k(G)$, called the k-inner product.

Thus we have shown that matrix coefficients of non-isomorphic representations are orthogonal for the k-inner product on $C_k(G)$.

Before going on, we can also exploit (4.7) in this situation: since $E^G = 0$, we derive

$$0 = \sum_{g \in G} \chi_E(g) \ ;$$

using the isomorphism

$$E \simeq E_1' \otimes E_2$$

which intertwines the action on E with $\tilde{\varrho}_1 \otimes \varrho_2$, we get

$$\chi_E(g) = \chi_{\pi_1}(g^{-1}) \chi_{\pi_2}(g),$$

and hence

$$[\chi_{\pi_2}, \chi_{\pi_1}] = \frac{1}{|G|} \sum_{g \in G} \chi_{\pi_1}(g^{-1}) \chi_{\pi_2}(g) = 0, \tag{4.10}$$

i.e., the characters of distinct irreducible representations are also orthogonal.

Now we consider the case where the representations π_1 and π_2 are equal; we then denote $\pi = \pi_1 = \pi_2$, acting on the space E. Then the space $\mathrm{End}_k(E)^G$ is one-dimensional and consists of the scalar operators. The precise application of the projection on the invariant space $\mathrm{End}_k(E)^G$ will now require to identify the scalar which is obtained.

[2] If $\varphi_1 \neq 0$, pick $x \in G$ with $\varphi_1(x) \neq 0$ and let φ_2 be the characteristic function of x^{-1}: then $[\varphi_1, \varphi_2] \neq 0$.

But first we apply (4.7), which does not involve such computations: the argument leading to (4.10) still applies, with the sole change that the trace of the projector is now 1. Hence we get

$$[\chi_\pi, \chi_\pi] = \frac{1}{|G|}\sum_{g\in G}\chi_\pi(g)\chi_\pi(g^{-1}) = 1. \tag{4.11}$$

We can proceed as before with rank 1 linear maps, but in the present case we should also observe that, because we deal with $\mathrm{End}_k(E)$, there are some other obvious linear maps to apply the projection to, namely the endomorphisms $\varrho(h)$, for $h \in G$.

We obtain that, for some $\lambda(h) \in k$, we have

$$\frac{1}{|G|}\sum_{g\in G} g\cdot\pi(h) = \frac{1}{|G|}\sum_{g\in G}\pi(ghg^{-1}) = \lambda(h)\mathrm{Id}_E.$$

To determine $\lambda(h)$, we take the trace (this is a standard technique): we get

$$\chi_\pi(h) = \lambda(h)\dim(\pi)$$

for any $h \in G$.

We have to be careful before concluding, because if k had positive characteristic, it might conceivably be the case that $\dim(\pi) = 0$ in k. However, we can see that this is not the case by noting that the formula just derived would then say that the character of π is identically 0, which contradicts either the linear independence of irreducible characters (or simply the formula (4.11).) Hence we have:

Proposition 4.3.4. *Let G and k be as above. For any irreducible k-representation π of G and any $h \in G$, we have*

$$\frac{1}{|G|}\sum_{g\in G}\pi(ghg^{-1}) = \frac{\chi_\pi(h)}{\dim(\pi)}.$$

We now finally come back to rank 1 maps. For given $w \in E$ and $\lambda \in E'$ defining

$$\Phi\ :\ \begin{cases} E \to E \\ v \mapsto \langle\lambda, v\rangle w, \end{cases}$$

as before, the operator

$$\frac{1}{|G|}\sum_{g\in G} g\cdot\Phi$$

is now a multiplication by a scalar α. To determine the scalar in question, we compute the trace. Since

$$\mathrm{Tr}(g\cdot\Phi) = \mathrm{Tr}(\pi(g)\Phi\pi(g^{-1})) = \mathrm{Tr}(\Phi)$$

for all g (using (2.15)), we get

$$(\dim E)\alpha = \mathrm{Tr}(\alpha\mathrm{Id}_E) = \mathrm{Tr}(\Phi) = \langle\lambda, w\rangle$$

(the trace of the rank 1 map can be computed by taking a basis where w, the generator of the image, is the first basis vector).

Since we have already seen that $\dim(E)$ is invertible in k, we obtain therefore

$$\frac{1}{|G|}\sum_{g\in G} g\cdot\Phi = \frac{\langle\lambda, w\rangle}{\dim E},$$

and applying this at a vector $v \in E$, and applying further a linear form $\mu \in E'$ to the result, we get:

Corollary 4.3.5 (Orthogonality of matrix coefficients). *With notation as above, let π be an irreducible k-representation of G. Then for all vectors $v, w \in E$ and linear forms $\lambda, \mu \in E'$, we have*

$$[f_{v,\lambda}, f_{w,\mu}] = \frac{1}{|G|} \sum_{g \in G} \langle \lambda, g \cdot v \rangle \langle \mu, g^{-1} \cdot w \rangle = \frac{\langle \lambda, w \rangle \langle \mu, v \rangle}{\dim E}. \tag{4.12}$$

We also summarize the orthogonality for characters:

Corollary 4.3.6 (Orthogonality of characters). *With notation as above, for any two irreducible representations π and ϱ of G, we have*

$$[\chi_\pi, \chi_\varrho] = \frac{1}{|G|} \sum_{g \in G} \chi_\pi(g) \chi_\varrho(g^{-1}) = \begin{cases} 0 & \textit{if } \pi \textit{ is not isomorphic to } \varrho, \\ 1 & \textit{otherwise,} \end{cases} \tag{4.13}$$

or equivalently

$$\frac{1}{|G|} \sum_{g \in G} \chi_{\pi(g)} \chi_{\tilde{\varrho}}(g) = \begin{cases} 0 & \textit{if } \pi \textit{ is not isomorphic to } \varrho \\ 1 & \textit{otherwise.} \end{cases} \tag{4.14}$$

Remark 4.3.7 (Invertibility of dimensions of irreducible representations). We have seen that for k algebraically closed of characteristic p not dividing $|G|$, the dimension of any irreducible representation is invertible in k. In fact, much more is known:

Theorem 4.3.8 (Divisibility). *Let G be a finite group. For any algebraically closed field of characteristic not dividing $|G|$, the family of the dimensions of irreducible k-representations of G is the same, and these dimensions divide the order of $|G|$.*

We will explain later how to show that $\dim(E) \mid |G|$ when $k = \mathbf{C}$ (Proposition 4.7.8); showing that the dimensions of the irreducible representations are the same for any algebraically closed field of characteristic $p \nmid |G|$ is more delicate, since it involves the Brauer characters mentioned in Remark 4.2.6.

4.3.3. Decomposition of class functions. As an application of the previous sections, we now give a slightly different proof of Theorem 4.2.4. The motivation is that the characters of irreducible k-representations of G are linearly independent class functions on G. The k-subspace they span in $C_k(G)$ is a subspace of $c_k(G)$ and the number of distinct characters is therefore at most $\dim c_k(G)$, which is the number of conjugacy classes. Hence the equality – which is the claim of the theorem – amounts to the statement that the characters actually generate $c_k(G)$, i.e., that they form a basis of this space.

We now prove this fact directly (yet another argument is contained in the proof we give of the corresponding facts for compact groups in Theorem 5.5.1). Let $\varphi \in c_k(G)$ be a class function. Like any function on G it can be expanded, at least, into a linear combination of matrix coefficients. We will show that, in this decomposition, only "diagonal" coefficients appear, and that those diagonal ones are constant (for a given irreducible representation), and this means that in fact the linear combination in question is a combination of characters.

To be precise, we fix, for every distinct irreducible representation π, a basis $(e_i^{(\pi)})_i$ of the space E_π of the representation, and denote by $(\lambda_j^{(\pi)})_j$ the dual basis. Let

$$f_{i,j}^{(\pi)} \in C_k(G), \qquad 1 \leqslant i, j \leqslant \dim(\pi),$$

denote the corresponding matrix coefficients. Theorem 2.7.24 (which amounts to the isotypic decomposition of the regular representation) shows that there exist coefficients $\alpha_{i,j}^{(\pi)} \in k$ such that

$$\varphi = \sum_{\pi} \sum_{i,j} \alpha_{i,j}^{(\pi)} f_{i,j}^{(\pi)}. \tag{4.15}$$

Our claims are: (1) for any π, and any distinct indices $i \neq j$, the coefficient $\alpha_{i,j}^{(\pi)}$ is zero; (2) for any π, the diagonal coefficients $\alpha_{i,i}^{(\pi)}$ are constant as i varies. Given this, if α_π denotes this last common value, we get

$$\varphi = \sum_{\pi} \alpha_\pi \chi_\pi$$

by interpreting the character as a sum of diagonal matrix coefficients.

To prove that claim, we use the orthogonality of matrix coefficients: using the choice of a basis and its dual, Corollaries 4.3.2 and 4.3.5 show that

$$[f_{i,j}^{(\pi)}, f_{k,l}^{(\varrho)}] = \begin{cases} 0 & \text{if } \pi \neq \varrho, \text{ or } (i,j) \neq (l,k) \\ \dfrac{1}{\dim(\varrho)} & \text{if } \pi = \varrho, \ (i,j) = (l,k). \end{cases}$$

Hence, taking the inner product with some $f_{k,l}^{(\varrho)}$ on both sides of (4.15), we get

$$[\varphi, f_{k,l}^{(\varrho)}] = \sum_{\pi} \sum_{i,j} \alpha_{i,j}^{(\pi)} [f_{i,j}^{(\pi)}, f_{k,l}^{(\varrho)}] = \frac{\alpha_{l,k}^{(\varrho)}}{\dim(\varrho)}. \tag{4.16}$$

We now think of ϱ as fixed. If we remember that $f_{k,l}^{(\varrho)}$ is the (l,k)-th coefficient of the matrix representing $\varrho(g)$ in the basis $(e_l^{(\varrho)})_l$, we can reinterpret the left-hand side of this computation as the (l,k)-th coefficient of the matrix representing

$$A_\varphi = \frac{1}{|G|} \sum_{g \in G} \varphi(g) \varrho(g^{-1}) \in \mathrm{End}_k(\varrho).$$

Now, *because φ is a class function*, it follows that, in fact, A_φ is in $\mathrm{End}_G(\varrho)$. This will be presented in the context of the group algebra later, but it is easy to check: we have

$$\begin{aligned} A_\varphi(\varrho(h)v) &= \frac{1}{|G|} \sum_{g \in G} \varphi(g) \varrho(g^{-1}h) v \\ &= \frac{1}{|G|} \sum_{g \in G} \varphi(g) \varrho(h(h^{-1}g^{-1}h)) v \\ &= \frac{1}{|G|} \sum_{g \in G} \varphi(hgh^{-1}) \varrho(h) \varrho(g^{-1}) v = \varrho(h) A_\varphi(v). \end{aligned}$$

Consequently, Schur's Lemma ensures – once again! – that A_φ is a scalar matrix. In particular, its off-diagonal coefficients $[\varphi, f_{k,l}^{(\varrho)}]$ with $k \neq l$ are zero, and the diagonal ones are constant; translating in terms of the coefficients $\alpha_{l,k}^{(\varrho)}$ using (4.16), we obtain the claim concerning the latter.

4.3.4. Orthogonality for unitary representations. We consider in this short section the case $k = \mathbf{C}$. Then we can proceed with the same arguments as in the previous example, but using the self-duality of Hilbert spaces, we may use the rank 1 linear maps

$$H_1 \xrightarrow{\Phi} H_2$$

between two Hilbert spaces defined by

$$\Phi(v) = \langle v, v_1 \rangle v_2$$

where $v_1 \in H_1$ and $v_2 \in H_2$, and the bracket denotes the inner product between vectors of H_1. The same analysis leads, when H_1 and H_2 are not isomorphic, to the relation

$$\frac{1}{|G|} \sum_{g \in G} g \cdot \Phi = 0,$$

and spelling it out by applying this to a $v \in E_1$ and taking the inner product of the result with $w \in E_2$, we obtain[3]

$$\begin{aligned} \frac{1}{|G|} \sum_{g \in G} \langle g \cdot v, v_1 \rangle \langle g^{-1} \cdot v_2, w \rangle &= \frac{1}{|G|} \sum_{g \in G} \langle g \cdot v, v_1 \rangle \overline{\langle g \cdot w, v_2 \rangle} \\ &= 0. \end{aligned}$$

We interpret this in terms of the invariant inner product

$$\langle \varphi_1, \varphi_2 \rangle = \frac{1}{|G|} \sum_{g \in G} \varphi_1(g) \overline{\varphi_2(g)}$$

for which the regular representation on the space of complex-valued functions $C(G)$ is unitary: the formula says that

$$\langle \varphi_{v_1,v_2}, \varphi_{v_3,v_4} \rangle = 0$$

for any "unitary" matrix-coefficients of non-isomorphic irreducible unitary representations of G.

Similarly, if $E = E_1 = E_2$ carries the irreducible unitary representation π, the same argument as in the previous example leads to

$$\frac{1}{|G|} \sum_{g \in G} g \cdot \Phi = \frac{\langle v_2, v_1 \rangle}{\dim E} \in \mathrm{End}_{\mathbf{C}}(E).$$

Applying to v and taking inner product with $w \in E$, we get

$$\begin{aligned} \frac{1}{|G|} \sum_{g \in G} \langle g \cdot v, v_1 \rangle \langle g^{-1} \cdot v_2, w \rangle &= \frac{1}{|G|} \sum_{g \in G} \langle g \cdot v, v_1 \rangle \overline{\langle g \cdot w, v_2 \rangle} \\ &= \frac{\langle v_2, v_1 \rangle \langle v, w \rangle}{\dim E}, \end{aligned}$$

i.e., renaming the vectors, we have

$$\langle \varphi_{v_1,v_2}, \varphi_{v_3,v_4} \rangle = \frac{\langle v_1, v_3 \rangle \overline{\langle v_2, v_4 \rangle}}{\dim E}$$

for any $v_i \in E$. Hence:

[3] Changing again g into g^{-1}.

Corollary 4.3.9 (Orthonormality of unitary matrix coefficients). *Let G be a finite group, $\pi : G \to \mathrm{U}(E)$ an irreducible unitary representation of G. For any orthonormal basis (e_i) of E, the normalized unitary matrix coefficients*

$$\varphi_{i,j} : x \mapsto \sqrt{\dim(E)}\langle \pi(x)e_i, e_j\rangle$$

are orthonormal in $C(G)$, with respect to the invariant inner product.

Indeed, we get

$$\langle \varphi_{i,j}, \varphi_{k,l}\rangle = \dim(E)\frac{\langle e_i, e_k\rangle\overline{\langle e_j, e_l\rangle}}{\dim E} = \begin{cases} 1 & \text{if } i = k \text{ and } j = l, \\ 0 & \text{otherwise.}\end{cases}$$

Similarly, for the characters themselves, we obtain the fundamental orthonormality of characters in the unitary case:

Corollary 4.3.10 (Orthonormality of characters). *Let G be a finite group and π, ϱ two irreducible unitary representations of G. We then have*

$$(4.17) \qquad \langle \chi_\varrho, \chi_\pi\rangle = \frac{1}{|G|}\sum_{g\in G}\chi_\pi(g)\overline{\chi_\varrho(g)} = \begin{cases} 0 & \textit{if } \pi \textit{ is not isomorphic to } \varrho, \\ 1 & \textit{otherwise.}\end{cases}$$

Hence the characters of irreducible unitary representations of G form an orthonormal basis of the space $c_k(G)$ of class functions on G with respect to the invariant inner product.

The last part is due to the fact that we know that the characters of irreducible unitary representations form an orthonormal family in $c_k(G)$, and that there are as many of them as there are conjugacy classes, i.e., as many as the dimension of $c_k(G)$, so that they must form an orthonormal basis.

4.3.5. Multiplicities. A crucial consequence of the orthogonality of characters is a formula for the multiplicities of irreducible representations in a given representation, at least in characteristic 0.

Proposition 4.3.11 (Multiplicities formula). *Let G be a finite group, and let k be an algebraically closed field of characteristic 0. For any finite-dimensional k-representation*

$$\varrho : G \longrightarrow \mathrm{GL}(E),$$

and for any irreducible k-representation π of G, the multiplicity $n_\pi(\varrho)$ of π in ϱ is given by

$$n_\pi(\varrho) = \dim \mathrm{Hom}_G(\pi, \varrho) = [\chi_\varrho, \chi_\pi] = \frac{1}{|G|}\sum_{g\in G}\chi_\varrho(g)\chi_\pi(g^{-1}).$$

If $k = \mathbf{C}$, then we can also write

$$n_\pi(\varrho) = \langle \chi_\pi, \chi_\varrho\rangle = \frac{1}{|G|}\sum_{g\in G}\chi_\varrho(g)\overline{\chi_\pi(g)}.$$

Note that we also have $n_\pi(\varrho) = \dim \mathrm{Hom}_G(\varrho, \pi)$ by the symmetry between irreducible quotient representations and subrepresentations, valid for a semisimple representation (Corollary 2.7.17).

Proof. Since ϱ is semisimple, we know that its character is given by

$$\chi_\varrho = \sum_\pi n_\pi(\varrho)\chi_\pi$$

in terms of the multiplicities and the various irreducible representations. Then we get by orthogonality

$$[\chi_\varrho, \chi_\pi] = n_\pi(\varrho)[\chi_\pi, \chi_\pi] = n_\pi(\varrho).$$

This is an equality in the field k, but since k has characteristic zero, it is also one in $\mathbf{Z}$ (in particular the left-hand side is an integer). □

More generally, we can extend this multiplicity formula by linearity:

PROPOSITION 4.3.12. *Let G be a finite group, and k an algebraically closed field of characteristic* 0. *For $i = 1$, 2, let*

$$\varrho_i : G \longrightarrow \mathrm{GL}(E_i)$$

be a finite-dimensional k-representation of G. Then we have

$$[\chi_{\varrho_1}, \chi_{\varrho_2}] = \dim \mathrm{Hom}_G(\varrho_1, \varrho_2) = \dim \mathrm{Hom}_G(\varrho_2, \varrho_1).$$

If $k = \mathbf{C}$, we have

$$\langle \chi_{\varrho_1}, \chi_{\varrho_2} \rangle = \dim \mathrm{Hom}_G(\varrho_1, \varrho_2) = \dim \mathrm{Hom}_G(\varrho_2, \varrho_1).$$

REMARK 4.3.13. It is customary to use (especially when $k = \mathbf{C}$) the shorthand notation

$$\langle \varrho_1, \varrho_2 \rangle = \langle \chi_{\varrho_1}, \chi_{\varrho_2} \rangle$$

for two representations ϱ_1 and ϱ_2 of G. We will do so to simplify notation, indicating sometimes the underlying group by writing $\langle \varrho_1, \varrho_2 \rangle_G$.

The multiplicity formula also leads to the following facts which are very useful when attempting to decompose a representation, when one doesn't know a priori all the irreducible representations of G. Indeed, this leads to a very convenient "numerical" criterion for irreducibility:

COROLLARY 4.3.14 (Irreducibility criterion). *Let G be a finite group, and let k be an algebraically closed field of characteristic* 0. *For any finite-dimensional k-representation ϱ of G, we have*

$$[\chi_\varrho, \chi_\varrho] = \sum_\pi n_\pi(\varrho)^2$$

where π runs over irreducible k-representations of G up to isomorphism, or, if $k = \mathbf{C}$, we have the formula

$$\langle \chi_\varrho, \chi_\varrho \rangle = \sum_\pi n_\pi(\varrho)^2$$

for the "squared norm" of the character of ϱ.

In particular, ϱ is irreducible if and only if

$$[\chi_\varrho, \chi_\varrho] = 1,$$

and if $k = \mathbf{C}$, if and only if

$$\langle \chi_\varrho, \chi_\varrho \rangle = \frac{1}{|G|} \sum_{g \in G} |\chi_\varrho(g)|^2 = 1.$$

PROOF. By linearity and orthogonality

$$[\chi_\varrho, \chi_\varrho] = \sum_{\pi_1, \pi_2} n_{\pi_1}(\varrho) n_{\pi_2}(\varrho) [\chi_{\pi_1}, \chi_{\pi_2}] = \sum_\pi n_\pi(\varrho)^2,$$

and similarly for $k = \mathbf{C}$. And if this is equal to 1, as an equality in $\mathbf{Z}$, the only possibility is that one of the multiplicities $n_\pi(\varrho)$ be equal to 1, and all the others are 0, which means $\varrho \simeq \pi$ is irreducible. □

Exercise 4.3.15 (Product groups). Let $G = G_1 \times G_2$ where G_1 and G_2 are finite groups. Use the irreducibility criterion to prove Proposition 2.3.17 directly for complex representations: all irreducible complex representations of G are of the form $\pi_1 \boxtimes \pi_2$ for some (unique) irreducible representations π_i of G_i.

Example 4.3.16 (Permutation representations). Suppose we have a complex representation ϱ of G with $\langle \chi_\varrho, \chi_\varrho \rangle = 2$. Then ϱ is necessarily a direct sum of two non-isomorphic irreducible subspaces, since 2 can only be written as $1^2 + 1^2$ as the sum of positive squares of integers.

A well-know source of examples of this is given by certain permutation representations (Section 2.6.2). Consider an action of G on a finite set X, and the associated permutation representation ϱ on the space E_X with basis vectors (e_x), so that

$$\varrho(g)e_x = e_{g\cdot x}.$$

The character of ϱ is given in Example 2.7.36: we have

$$\chi_\varrho(g) = |\{x \in X \mid g \cdot x = x\}|.$$

We first deduce from this that

$$\langle \chi_\varrho, \mathbf{1} \rangle = \frac{1}{|G|} \sum_{g\in G} \sum_{\substack{x\in X\\ g\cdot x = x}} 1 = \frac{1}{|G|} \sum_{x\in X} |G_x|$$

where $G_x = \{g \in G \mid g \cdot x = x\}$ is the stabilizer of x in G.

The order of this subgroup depends only on the orbit of x: if $y = g \cdot x$, we have

$$G_y = gG_xg^{-1}.$$

Hence, summing over the orbits, we get

$$\langle \chi_\varrho, \mathbf{1} \rangle = \sum_{o\in G\backslash X} \frac{|G_o||o|}{|G|} = |G\backslash X|, \tag{4.18}$$

the number of orbits (we used the standard bijection $G_o\backslash G \longrightarrow o$ induced by mapping $g \in G$ to $g \cdot x_0$ for some fixed $x_0 \in o$).

We assume now that there is a single orbit, i.e., that the action of G on X is transitive (otherwise, ϱ already contains at least two copies of the trivial representation). Then, since the character of ϱ is real-valued, we have

$$\begin{aligned}
\langle \varrho, \varrho \rangle &= \frac{1}{|G|} \sum_{g\in G} \Bigl(\sum_{\substack{x\in X\\ g\cdot x = x}} 1 \Bigr)^2 \\
&= \frac{1}{|G|} \sum_{x,y\in X} \sum_{g\in G_x\cap G_y} 1 \\
&= 1 + \frac{1}{|G|} \sum_{x\neq y} \sum_{g\in G_x\cap G_y} 1 \\
&= 1 + \frac{1}{|G|} \sum_{g\in G} |\{(x,y) \in X \times X \mid x \neq y \text{ and } g\cdot(x,y) = (x,y)\}|
\end{aligned}$$

and we recognize, from the character formula for a permutation representation, that this means

$$\langle \varrho, \varrho \rangle = 1 + \langle \varrho^{(2)}, \mathbf{1} \rangle,$$

where $\varrho^{(2)}$ is the permutation representation corresponding to the natural action of G on the set

$$Y = \{(x, y) \in X \mid x \neq y\}$$

(recall that $g \cdot x = g \cdot y$ implies that $x = y$, so the action of G on $X \times X$ leaves Y invariant.) By (4.18), we see that we have $\langle \varrho, \varrho \rangle = 2$ if (and, in fact, only if) this action on Y is *transitive*. This, by definition, is saying that the original action was *doubly transitive*: not only can G bring any element $x \in X$ to any other (transitivity), but a single element can simultaneously bring any x to any x', and any y to any y', provided the conditions $x \neq y$ and $x' \neq y'$ are satisfied.

Thus:

Proposition 4.3.17. *Let G be a finite group acting doubly transitively on a finite set X. Then the representation of G on the space*

$$E = \{\sum_{x \in X} \lambda_x e_x \mid \sum_{x \in X} \lambda_x = 0\}$$

induced by $\varrho(g)e_x = e_{gx}$ is an irreducible complex representation of G of dimension $|X|-1$, with character

$$\chi_\varrho(g) = |\{x \in X \mid g \cdot x = x\}| - 1.$$

Indeed, the subspace E is stable under the action of G (it is the orthogonal of the space E_X^G for the natural inner product on E_X such that (e_x) is an orthonormal basis.)

For a concrete example, consider $G = \mathfrak{S}_n$ acting on $X = \{1, \ldots, n\}$ by permutations. If $n \geqslant 2$, this action is doubly transitive (as the reader should make sure to check, if needed!), and this means that the representation of $\mathfrak{S}_n$ on the hyperplane

$$E_n = \{(x_i) \in \mathbf{C}^n \mid \sum_i x_i = 0\} \tag{4.19}$$

is irreducible of dimension $n - 1$.

Remark 4.3.18. A warning about the irreducibility criterion: it only applies if one knows that χ_ϱ is, indeed, the character of a representation of G. There are many class functions φ with squared norm 1 which are not characters, for instance

$$\varphi = \frac{3\chi_{\pi_1} + 4\chi_{\pi_2}}{5}$$

if π_1 and π_2 are non-isomorphic irreducible representations. If π_1 and π_2 have dimension divisible by 5, the non-integrality might not be obvious from looking simply at the character values!

However, note that if $\varphi \in R(G)$ is a *virtual character* (over $\mathbf{C}$, say), i.e., $\varphi = \chi_{\varrho_1} - \chi_{\varrho_2}$ for some actual complex representations ϱ_1 and ϱ_2, the condition

$$\langle \varphi, \varphi \rangle = 1$$

means that either ϱ_1 or ϱ_2 is irreducible and the other zero, or in other words, either φ or $-\varphi$ is a character of an irreducible representation of G. Indeed, we can write

$$\varphi = \sum_\pi n_\pi \chi_\pi$$

as a combination of irreducible characters with *integral* coefficients n_π, and we have again

$$\langle \varphi, \varphi \rangle = \sum_\pi n_\pi^2$$

so one, and only one, of the n_π is equal to ± 1, and the others are 0.

Example 4.3.19 (Frobenius reciprocity). From the general multiplicity formula, we get a "numerical" version of Frobenius reciprocity for induced representations (Proposition 2.3.6): given a subgroup H of a finite group G, a (complex, say) representation ϱ_1 of G and a representation ϱ_2 of H, we have[4]

$$\langle \varrho_1, \mathrm{Ind}_H^G(\varrho_2)\rangle_G = \langle \mathrm{Res}_H^G \varrho_1, \varrho_2\rangle_H.$$

Note that, by symmetry, we also have

$$\langle \mathrm{Ind}_H^G(\varrho_2), \varrho_1\rangle_G = \langle \varrho_2, \mathrm{Res}_H^G \varrho_1\rangle_H,$$

(something which is not universally true in the generality in which we defined induced representations.)

This numerical form of Frobenius reciprocity can easily be checked directly, as an identity between characters: denoting $\chi_i = \chi_{\varrho_i}$, we find using (2.45) that we have

$$\begin{aligned}
\langle \varrho_1, \mathrm{Ind}_H^G(\varrho_2)\rangle_G &= \frac{1}{|G|}\sum_{g\in G}\chi_1(g)\overline{\sum_{\substack{s\in H\backslash G\\ sgs^{-1}\in H}}\chi_2(sgs^{-1})}\\
&= \frac{1}{|G|}\sum_{s\in H\backslash G}\sum_{g\in s^{-1}Hs}\chi_1(g)\overline{\chi_2(sgs^{-1})}\\
&= \frac{1}{|G|}\sum_{s\in H\backslash G}\sum_{h\in H}\chi_1(s^{-1}hs)\overline{\chi_2(h)}\\
&= \frac{|H\backslash G|}{|G|}\sum_{h\in H}\chi_1(h)\overline{\chi_2(h)} = \langle \mathrm{Res}_H^G(\varrho_1), \varrho_2\rangle_H.
\end{aligned}$$

For instance, if one thinks that ϱ_1 is an irreducible representation of G, and ϱ_2 is one of H, Frobenius reciprocity says that "the multiplicity of ϱ_1 in the representation induced from ϱ_2 is the same as the multiplicity of ϱ_2 in the restriction of ϱ_1 to H."

Here is an example of application: for a finite group G, we denote by $A(G)$ the maximal dimension of an irreducible complex representation of G. So, for instance, $A(G) = 1$ characterizes finite abelian groups. More generally, $A(G)$ can be seen to be some measure of the complexity of G.

Proposition 4.3.20. *For any finite group G and subgroup $H \subset G$, we have $A(H) \leqslant A(G)$.*

Proof. To see this, pick an irreducible representation π of H such that $\dim \pi = A(H)$. Now consider the induced representation $\varrho = \mathrm{Ind}_H^G(\pi)$. It may or may not be irreducible; in any case, let τ be any irreducible component of ϱ; then we have

$$1 \leqslant \langle \tau, \varrho\rangle = \langle \tau, \mathrm{Ind}_H^G(\pi)\rangle = \langle \mathrm{Res}_H^G(\tau), \pi\rangle_H$$

by Frobenius reciprocity. This means that π occurs with multiplicity at least 1 in the restriction of τ. This implies that necessarily $\dim \tau = \dim \mathrm{Res}(\tau) \geqslant \dim \pi$. Thus τ is an irreducible representation of G of dimension at least $A(H)$, i.e., we have $A(G) \geqslant A(H)$. □

[4] We use the notation of Remark 4.3.13.

EXERCISE 4.3.21. For a finite group G and a real number $p \geqslant 0$, let

$$A_p(G) = \sum_{\pi} (\dim \pi)^p.$$

If $p \geqslant 1$, show that for any subgroup $H \subset G$, we have $A_p(H) \leqslant A_p(G)$.

Here is a last, very cute, application of the multiplicity formula:

PROPOSITION 4.3.22 (Where to find irreducible representations?). *Let G be a finite group and let $\varrho : G \longrightarrow \mathrm{GL}(E)$ be any finite-dimensional faithful complete representation. Then any irreducible representation $\pi \in \hat{G}$ can be found as a subrepresentation of a tensor power $\varrho \otimes \cdots \otimes \varrho$, with k factors, for some $k \geqslant 1$.*

PROOF. Fix $\pi \in \hat{G}$, and define $m_k \geqslant 0$ as the multiplicity of π in $\varrho^{\otimes k}$ for $k \geqslant 0$ (with the convention that the 0-th tensor power is the trivial representation), in other words

$$m_k = \langle \varrho^{\otimes k}, \pi \rangle$$

for $k \geqslant 0$. The goal is therefore to show that this multiplicity m_k is non-zero for some $k \geqslant 0$. The clever idea is to consider the generating series

$$\sum_{k \geqslant 0} m_k X^k \in \mathbf{Z}[[X]],$$

and show that it can not be zero. For this we write

$$\langle \varrho^{\otimes k}, \pi \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_{\varrho}(g)^k \overline{\chi_{\pi}(g)},$$

and compute the power series by exchanging the two sums:

$$\sum_{k \geqslant 0} m_k X^k = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_{\pi}(g)} \sum_{k \geqslant 0} \chi_{\varrho}(g)^k X^k = \frac{1}{|G|} \sum_{g \in G} \frac{\overline{\chi_{\pi}(g)}}{1 - \chi_{\varrho}(g) X}.$$

This doesn't look like the 0 power series, but there might be cancellations in the sum. However, we haven't used the assumption that ϱ is faithful, and there is the cunning trick: the point $1/\chi_{\varrho}(1) = 1/\dim(\varrho)$ is a pole of the term corresponding to $g = 1$, and it *can not* be cancelled because $\chi_{\varrho}(g) = \dim(\varrho)$ if and only if $g \in \ker \varrho = 1$ (this is the easy Proposition 4.6.4 below; the point is that the character values are traces of unitary matrices, hence sum of $\dim \varrho$ complex numbers of modulus 1.) So it follows that the power series is non-zero, which certainly means that m_k is not always 0. □

4.3.6. Isotypic projectors. We now come back to the problem of determining the projectors on all the isotypic components of a representation, not only the invariant subspace. In the language of the group algebra, Proposition 4.3.1 means that the *single* element

$$e = \frac{1}{|G|} \sum_{g \in G} g \in k(G)$$

of the group algebra has the property that its action on *any* representation of G gives "universally" the space of invariants. Since E^G is the same as the isotypic component of E with respect to the trivial representation, it is natural to ask for similar elements for the other irreducible representations of G. These exist indeed, and they are also remarkably simple: they are given by the characters.

Proposition 4.3.23 (Projectors on isotypic components). *Let G be a finite group, k an algebraically closed field of characteristic $p \nmid |G|$. For any k-representation*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of G, and for any irreducible k-representation π of G, the element

$$e_\pi = \frac{\dim(\pi)}{|G|} \sum_{g \in G} \chi_\pi(g^{-1}) g \in k(G)$$

acts on E as a homomorphism in $\mathrm{Hom}_G(E, E)$ and is a projector onto the isotypic component $M(\pi) \subset E$.

In other words, the linear map

$$\begin{cases} E \longrightarrow E \\ v \mapsto \dfrac{\dim \pi}{|G|} \displaystyle\sum_{g \in G} \chi_\pi(g^{-1}) \varrho(g) v \end{cases} \tag{4.20}$$

is a G-homomorphism, and is a projection onto $M(\pi)$.

For $k = \mathbf{C}$, if we think of unitary representations, we get:

Proposition 4.3.24 (Orthogonal projectors on isotypic components). *Let G be a finite group and let $\varrho : G \longrightarrow \mathrm{U}(H)$ be a unitary representation of G. For any irreducible unitary representation π of G, the element*

$$e_\pi = \frac{\dim(\pi)}{|G|} \sum_{g \in G} \overline{\chi_\pi(g)} g \in \mathbf{C}(G)$$

acts on E as a homomorphism in $\mathrm{Hom}_G(E, E)$ and is the orthogonal *projector onto the isotypic component $M(\pi) \subset E$.*

We will explain how one can find this formula for e_π, instead of merely checking its properties. Indeed, this leads to additional insights. The point is that, for a given irreducible representation π, the family of projections to the π-isotypic component, which maps all others to 0, gives for every representation $\varrho : G \longrightarrow \mathrm{GL}(E)$ of G a linear map

$$\varepsilon_\varrho : E \longrightarrow E,$$

in a "functorial" manner, in the sense described in Exercise 3.1.5: for any representation τ on F and any G-homomorphism

$$E \xrightarrow{\Phi} F,$$

we have

$$\varepsilon_\tau \circ \Phi = \Phi \circ \varepsilon_\varrho.$$

Exercise 4.3.25. Check this fact (this is because Φ sends the isotypic components $M_E(\tau)$ to $M_F(\tau)$, for any irreducible representation τ (Proposition 2.7.7, (3)).

The outcome of Exercise 3.1.5 is that the source of a "universal" linear map on all representations can only be the action of some *fixed* element a of the group algebra; even if you did not solve this exercise, it should be intuitively reasonable that this is the only obvious source of such maps. Thus, we know a priori that there *is* a formula for the projector. We only need to find it.

The projectors are not just linear maps, but also intertwiners; according to the last part of Exercise 3.1.5, this corresponds to an element a of the group algebra $k(G)$ which is in its center $Z(k(G))$. (This is because a gives rise to G-homomorphism if and only if a satisfies

$$g \cdot a = a \cdot g \in k(G)$$

for all $g \in G$, which is equivalent with $a \in Z(k(G))$ because G generates $k(G)$ as a ring.)

Remark 4.3.26. If we write

$$a = \sum_{x \in G} \alpha_x x, \qquad \alpha_x \in k,$$

the condition that a belong to the center becomes

$$\alpha_{x^{-1}g} = \alpha_{gx^{-1}}, \qquad \text{for all } x \text{ and } g,$$

or, in other words, the function

$$x \mapsto \alpha_x$$

must be a *class function.*

Now we assume that $a \in Z(k(G))$, so that a acts as a G-morphism on every representation of G. In particular, the action of a on an irreducible representation π must be given by multiplication by some scalar $\omega_\pi(a) \in k$, according to Schur's Lemma. Because this is "universal", we see that the element giving the projection on $M(\pi)$ is the element $a \in k(G)$ such that $\omega_\pi(a) = 1$, and $\omega_\tau(a) = 0$ for all other (non-isomorphic) irreducible k-representations τ of G – indeed, if a has this property, it follows that for a given representation of G on E, a acts as identity on all subrepresentations of E isomorphic to π, i.e., on $M(\pi)$, and also a that annihilates all other isotypic components. This is exactly the desired behavior.

To determine a exactly, we observe that we can compute $\omega_\tau(a)$, as a function of the coefficients α_x of a and of the irreducible representation τ, by taking the trace: from

$$\omega_\tau(a)\mathrm{Id}_\tau = \sum_{x \in G} \alpha_x \tau(x),$$

we get

$$\omega_\tau(a) \dim \tau = \sum_{x \in G} \alpha_x \operatorname{Tr}(\tau(x)) = \sum_{x \in G} \alpha_x \chi_\tau(x).$$

Hence we are looking for coefficients α_x such that

$$\sum_{x \in G} \alpha_x \chi_\pi(x) = \dim \pi$$

(the case $\tau = \pi$) and

$$\sum_{x \in G} \alpha_x \chi_\tau(x) = 0,$$

if τ is an irreducible representation non-isomorphic to π. But the orthogonality of characters (4.13) precisely says that

$$\alpha_x = \dim(\pi)\chi_\pi(x^{-1})$$

satisfies these conditions, and when $k = \mathbf{C}$ and we have unitary representations, this becomes

$$\alpha_x = \dim(\pi)\overline{\chi_\pi(x)}.$$

(see also Corollary 4.3.10.) Thus Proposition 4.3.24 is proved.

Having obtained the formula for the projectors on isotypic components of any representation, there is one important example that should come to mind where we can (and should) apply this: the group algebra itself, when G acts on $k(G)$ by multiplication on the left. The special feature of $k(G)$ is its algebra structure, which also gives some extra structure to the isotypic components.

Let $I(\pi)$ be the π-isotypic component of $k(G)$. According to the above, the projection on $I(\pi)$ is given by $a \mapsto e_\pi a$, where

$$e_\pi = \frac{\dim(\pi)}{|G|} \sum_{g \in G} \chi(g^{-1}) g.$$

Taking $a = 1$, we deduce from this that $e_\pi \in I(\pi)$ in particular. This means for instance that

$$e_\pi^2 = e_\pi e_\pi = e_\pi, \tag{4.21}$$

and also (since other projections map $I(\pi)$ to 0) that

$$e_\varrho e_\pi = 0 \tag{4.22}$$

if ϱ is an irreducible representation not isomorphic to π. Note moreover that

$$1 = \sum_\pi e_\pi, \tag{4.23}$$

which is simply because of the isotypic decomposition

$$k(G) = \bigoplus_\pi I(\pi).$$

In any ring A, a family (e_i) of elements satisfying the relations (4.21), (4.22) and (4.23) is known as a "complete system of independent idempotents". Their meaning is the following:

Corollary 4.3.27 (Product decomposition of the group algebra). *Let G be a finite group and k an algebraically closed field of characteristic not dividing $|G|$. Then the subspaces $I(\pi)$, where π runs over irreducible k-representations of G, are two-sided ideals in $k(G)$. Moreover, with $e_\pi \in I(\pi)$ as unit, $I(\pi)$ is a subalgebra of $k(G)$ isomorphic to the matrix algebra $\mathrm{End}_k(\pi)$, and we have a k-algebra isomorphism*

$$\begin{cases} k(G) & \xrightarrow{\sim} & \prod_\pi \mathrm{End}(\pi) \\ a & \mapsto & (\pi(e_\pi a))_\pi \end{cases} \tag{4.24}$$

Proof. The space $I(\pi)$ is the image of the projection given by multiplication by e_π, i.e., we have

$$I(\pi) = e_\pi k(G),$$

which is, a priori, a right-ideal in $k(G)$. But if we remember that e_π is also in the center $Z(k(G))$ of the group algebra, we deduce that $I(\pi) = k(G)e_\pi$ also, i.e., that $I(\pi)$ is a two-sided ideal.

In particular, like any two-sided ideal, $I(\pi)$ is stable under multiplication. Usually, $1 \notin I(\pi)$, so that 1 does not provide a unit. But, for any $a \in I(\pi)$, if we write $a = e_\pi a_1$, we find that

$$e_\pi a = e_\pi^2 a_1 = e_\pi a_1 = a$$

by (4.21), and similarly $ae_\pi = a$, so that e_π, which is in $I(\pi)$, is indeed a unit for this two-sided ideal!

The identity (4.23) means that, as algebras, we have

$$k(G) \simeq \prod_{\pi} I(\pi),$$

where any $a \in k(G)$ corresponds to $(e_\pi a)_\pi$. Thus there only remains to prove that the map

$$\begin{cases} I(\pi) & \longrightarrow & \mathrm{End}_k(\pi) \\ a & \mapsto & \pi(a). \end{cases}$$

is an algebra isomorphism.

This is certainly an algebra homomorphism (the unit e_π maps to the identity in $\mathrm{End}_k(\pi)$, since – by the above – the action of e_π is the projector on $M(\pi)$, which is the identity for π itself.) It is surjective, by Burnside's irreducibility criterion (the latter says, more precisely, that the image of all of $k(G)$ is $\mathrm{End}_k(\pi)$, but the other isotypic components map to 0.) We can show that it is an isomorphism either by dimension count (since the representation of G on $k(G)$ is, for a finite group, isomorphic to that on $C_k(G)$, the isotypic component $I(\pi)$ has the same dimension as the space of matrix coefficients of π, namely $(\dim \pi)^2$), or by proving injectivity directly: if $a \in I(\pi)$ satisfies $\pi(a) = 0$, it follows that the action of a on every π-isotypic component of every representation is also zero; if we take the special case of $k(G)$ itself, this means in particular that $ae_\pi = 0$. However, $ae_\pi = e_\pi a = a$ if $a \in I(\pi)$, and thus $a = 0$. □

Remark 4.3.28. This result also leads to Theorem 4.2.4: the center $Z(k(G))$ of $k(G)$ has dimension equal to the number of conjugacy classes of G (since it is the space of all elements

$$a = \sum_{g \in G} \alpha_g g,$$

with $g \mapsto \alpha_g$ a class function, as we observed before), while from (4.24), we have

$$Z(k(G)) \simeq \prod_{\pi} Z(\mathrm{End}_k(\pi)) = \prod_{\pi} k\mathrm{Id}_\pi$$

(since any endomorphism ring has one-dimensional center spanned by the identity). Thus the dimension of $Z(k(G))$ is also equal to the number of irreducible k-representations of G.

Exercise 4.3.29 (How big can a cyclic representation be?). Let G be a finite group and k a field of characteristic not dividing $|G|$. Let $\varrho : G \longrightarrow \mathrm{GL}(E)$ be a finite-dimensional k-representation of G, and π an irreducible k-representation.

(1) For $v \in M_E(\pi) \subset E$, show that the subrepresentation F_v of E generated by v (which is a cyclic representation, see Remark 2.2.6) is the direct sum of at most $\dim(\pi)$ copies of π. [Hint: Show that F_v is isomorphic to a quotient of the π-isotypic component of $k(G)$.]

(2) Show that this can not be improved (i.e., it is possible, for some ϱ and v, that F_v is the direct sum of exactly $\dim(\pi)$ copies of π.)

(3) If you solved (1) using the group algebra $k(G)$, try to do it without (see [**34**, Ex. 2.10] if needed).

In the argument leading to the projector formula, we have also proved the following useful result:

Proposition 4.3.30 (Action of the center of the group ring). *Let G be a finite group and k an algebraically closed field of characteristic not dividing $|G|$. For any irreducible k-representation ϱ of G, there is an associated algebra homomorphism*

$$\omega_\varrho : \begin{cases} Z(k(G)) & \longrightarrow & k \\ a & \mapsto & \varrho(a), \end{cases}$$

i.e.,

$$\varrho(a) = \omega_\varrho(a)\mathrm{Id}. \tag{4.25}$$

This is given by

$$\omega_\varrho\Bigl(\sum_{g\in G} \alpha_g g\Bigr) = \frac{1}{\dim(\varrho)} \sum_{g\in G} \alpha_g \chi_\varrho(g). \tag{4.26}$$

The last formula is obtained, as usual, by taking the trace on both sides of (4.25). Note the following special case: if $c \subset G$ is a conjugacy class, the element

$$a_c = \sum_{g\in c} g$$

is in the center of the group algebra, and we get

$$\omega_\varrho(a_c) = \frac{|c|\chi_\varrho(c)}{\dim \varrho}. \tag{4.27}$$

This can be used to show how to compute in principle all characters of irreducible representations of G: see Proposition 4.6.2 below.

4.4. Harmonic analysis on finite groups

The terminology "harmonic analysis" refers roughly to the use of specific orthonormal bases of a Hilbert space to analyze its elements, and in particular to cases of function spaces. In the case of finite groups, there are two main examples, which are related: (1) either one considers the space $c(G)$ of complex-valued class functions, and the orthonormal basis of irreducible characters; (2) or one considers the full space $C(G)$ of functions on the group, and a basis of matrix coefficients. The second case is often less easy to handle, because matrix coefficients are not entirely canonical objects. This explains also why the first case is worth considering separately, and not simply as a corollary of the theory of matrix coefficients.

Given a class function $f \in c(G)$, we have

$$f = \sum_{\varrho\in\hat{G}} \langle f, \chi_\varrho\rangle \chi_\varrho$$

where $\hat{G}$ denotes the set of irreducible complex representations of G, up to isomorphism. It is worth isolating the contribution of the trivial representation $\mathbf{1} \in \hat{G}$, which is the constant function with value

$$\langle f, 1\rangle = \frac{1}{|G|} \sum_{g\in G} f(g)$$

i.e., the average value of f on G. It is characteristic of harmonic analysis to decompose f in such a way that its "average" behavior is clearly separated from the fluctuations around it, which are given by the sum of the contributions of non-trivial characters.

We now write down "explicitly" what is the outcome of this decomposition when f is taken to be especially simple. Precisely, fix a conjugacy class $c \subset G$, and let f_c be its characteristic function, which is a class function. We then obtain:

Corollary 4.4.1 (Decomposition of characteristic functions of conjugacy classes). *Let g and $h \in G$. We have*

$$\sum_{\varrho\in\hat{G}} \chi_\varrho(h)\overline{\chi_\varrho(g)} = \begin{cases} \dfrac{|G|}{|g^\sharp|} & \text{if } g \text{ is conjugate to } h, \\ 0 & \text{otherwise,} \end{cases} \tag{4.28}$$

where $g^\sharp$ is the conjugacy class of g.

This corollary is often called the "second orthogonality formula", and is usually proved by observing that the transpose of a unitary matrix (namely, the character table of G) is still unitary. Note that in the "diagonal" case, the value

$$\frac{|G|}{|g^\sharp|}$$

is also equal to $|C_G(g)|$, the size of the centralizer of g in G.

Proof. As indicated, we expand f_c in terms of characters:

$$f_c = \sum_{\varrho\in\hat{G}} \langle f_c, \chi_\varrho\rangle \chi_\varrho$$

and we remark that, by definition, we have

$$\langle f_c, \chi_\varrho\rangle = \frac{1}{|G|}\sum_{g\in G} f_c(g)\overline{\chi_\varrho(g)} = \frac{|c|}{|G|}\overline{\chi_\varrho(h)}$$

since f_c is 1 on the conjugacy class c and 0 elsewhere. □

Remark 4.4.2 (The space of conjugacy classes). The space $c(G)$ of class functions can be identified with the space $C(G^\sharp)$ of complex-valued functions on the set $G^\sharp$ of *conjugacy classes* in G, since a class function is constant on each conjugacy class. It is often useful to think in these terms. However, one must be careful that the Hilbert space inner product on $C(G^\sharp)$ coming from this identification (i.e., the inner product such that the identification is an isometry) is *not* the inner product

$$\frac{1}{|G^\sharp|}\sum_{c\in G^\sharp} f_1(c)\overline{f_2(c)}$$

that might seem most natural on a finite set. Instead, we have

$$\langle f_1, f_2\rangle = \frac{1}{G}\sum_{c\in G^\sharp} |c| f_1(c)\overline{f_2(c)}$$

for any functions

$$f_1,\ f_2\ :\ G^\sharp \longrightarrow \mathbf{C}.$$

This means that each conjugacy class carries a weight which is proportional to its size as a subset of G, instead of being uniformly distributed.

Harmonic analysis often involves using the expansion of a characteristic function in order to replace a condition of the type "g is in such and such subset X of G" by its expansion in terms of some orthonormal basis, so that one, for instance, write

$$\sum_{x\in X} f(x) = \sum_{x\in G} f(x)\mathbf{1}_X(x) = \sum_{\text{basis }(\varphi_i)} \langle \mathbf{1}_X, \varphi_i\rangle \sum_{x\in G} f(x)\varphi_i(x)$$

(where $\mathbf{1}_X$ is the characteristic function of X). Furthermore, it is usually the case that the constant function 1 is part of the orthonormal basis (this is the case for characters as for matrix coefficients), in which case the corresponding term is

$$\langle \mathbf{1}_X, 1\rangle \sum_{x\in G} f(x) = \frac{|X|}{|G|}\sum_{x\in X} f(x),$$

which can be interpreted as the term that would arise from a heuristic argument, where $|X|/|G|$ is seen as the probability that some element of G is in X.

We present a first interesting illustration here, which is due to Frobenius, and another one is found in Section 4.7.1 later on.

Proposition 4.4.3 (Detecting commutators). *Let G be a finite group. The number $N(g)$ of $(x,y)\in G\times G$ such that $g=[x,y]=xyx^{-1}y^{-1}$ is equal to*

$$|G|\sum_{\pi\in\hat{G}} \frac{\chi_\pi(g)}{\chi_\pi(1)}, \tag{4.29}$$

and in particular g is a commutator if and only if

$$\sum_{\pi\in\hat{G}} \frac{\chi_\pi(g)}{\chi_\pi(1)} \neq 0.$$

Proof. This will be a bit of a roller-coaster, and we won't try to motivate the arguments... The first idea is to compute instead $N^\sharp(g)$, the number of $(x,y)\in G\times G$ such that $[x,y]$ is *conjugate* to G. The point is that

$$N^\sharp(g) = \sum_{h\in g^\sharp} N(h) = |g^\sharp|N(g) \tag{4.30}$$

because the representations of conjugate elements as commutators are naturally in bijection:

$$[x,y]=g \text{ if and only if } [zxz^{-1},zyz^{-1}] = zgz^{-1}\in g^\sharp,$$

so that one recovers easily $N(g)$ from $N(g^\sharp)$, while relaxing equality to conjugation allows us to detect the condition using characters instead of the full matrix coefficients.

Now we start by fixing x, and attempt to determine the number $n(x,g)$ of $y\in G$ such that g is conjugate to $[x,y]$, so that

$$N(g^\sharp) = \sum_{x\in G} n(x,g).$$

We compute $n(x,g)$ using characters: we have

$$n(x,g) = \sum_{y\in G} f_g([x,y]),$$

where f_g denotes the characteristic function of the conjugacy class of g. Using Corollary 4.4.1, and exchanging the order of the two sums, we get

$$n(x,g) = \frac{|g^\sharp|}{|G|} \sum_{\pi\in\hat{G}} \overline{\chi_\pi(g)} \sum_{y\in G} \chi_\pi([x,y]). \tag{4.31}$$

In order to go further, we consider the inner sum as $\beta(x,x^{-1})$ where β is a function of two variables:

$$\beta(a,b) = \sum_{y\in G} \chi_\pi(ayby^{-1}).$$

Note that β itself is a class function of a: we have

$$\begin{aligned}\beta(hah^{-1},b) = \sum_{y\in G} \chi_\pi(hah^{-1}yby^{-1}) &= \sum_{y\in G} \chi_\pi(hah^{-1}yby^{-1}) \\ &= \sum_{y\in G} \chi_\pi(h^{-1}yby^{-1}ha) = \sum_{y\in G} \chi_\pi(yby^{-1}a) = \beta(a,b)\end{aligned}$$

(after the change of variable $h^{-1}y \mapsto y$). We therefore attempt to expand β in terms of characters, with respect to the a-variable:

$$\beta(a,b) = \sum_{\varrho\in\hat{G}} \Big(\frac{1}{|G|}\sum_{h\in G} \beta(h,b)\overline{\chi_\varrho(h)}\Big)\chi_\varrho(a). \tag{4.32}$$

We are not going in circle, and we look at the inner coefficient:

$$\frac{1}{|G|}\sum_{h\in G}\beta(h,b)\overline{\chi_\varrho(h)} = \frac{1}{|G|}\sum_{h\in G}\sum_{y\in G}\chi_\pi(hyby^{-1})\overline{\chi_\varrho(h)} = \frac{1}{|G|}\sum_{y\in G}\sum_{h\in G}\chi_\pi(hyby^{-1})\overline{\chi_\varrho(h)}.$$

At last, the inner sum *here* can be recognized: the function

$$z \mapsto \frac{1}{|G|}\sum_{h\in G}\chi_\pi(hz)\overline{\chi_\varrho(h)} = \frac{1}{|G|}\sum_{h\in G}\chi_\pi(zh)\overline{\chi_\varrho(h)}$$

is simply

$$\frac{1}{|G|}\sum_{h\in G}\overline{\chi_\varrho(h)}\,\mathrm{reg}(h)\chi_\pi,$$

or in other words (Proposition 4.3.23), it is $1/\dim(\varrho)$ times the image of χ_π under the projection on the ϱ-isotypic component of the regular representation! Since χ_π is in the π-isotypic component, this projection is 0 except when $\varrho = \pi$, when it is equal to π. Hence

$$\frac{1}{|G|}\sum_{h\in G}\chi_\pi(hz)\overline{\chi_\varrho(h)} = \begin{cases}\frac{\chi_\pi(z)}{\dim(\pi)} & \text{if } \varrho\simeq\pi\\ 0 & \text{otherwise,}\end{cases}$$

and

$$\frac{1}{|G|}\sum_{h\in G}\beta(h,b)\overline{\chi_\varrho(h)} = \begin{cases}|G|\frac{\chi_\pi(yby^{-1})}{\dim(\pi)} & \text{if } \varrho\simeq\pi\\ 0 & \text{otherwise.}\end{cases}$$

Going back to (4.32), only the term $\varrho = \pi$ survives, and gives

$$\beta(a,b) = \frac{1}{\dim\pi}\chi_\pi(a)\chi_\pi(yby^{-1}) = \frac{|G|}{\dim\pi}\chi_\pi(a)\chi_\pi(b),$$

or in other words, we have the fairly nice formula

$$\chi_\pi(a)\chi_\pi(b) = \frac{(\dim \pi)}{|G|} \sum_{y\in G} \chi_\pi(ayby^{-1}). \tag{4.33}$$

We can now come back to $n(x,g)$, i.e., to (4.31): we have

$$n(x,g) = \frac{|g^\sharp|}{|G|} \sum_{\pi\in\hat{G}} \overline{\chi_\pi(g)}\beta(x,x^{-1}) = |g^\sharp| \sum_{\pi\in\hat{G}} \frac{\overline{\chi_\pi(g)}|\chi_\pi(x)|^2}{\dim \pi}.$$

Summing over x, we get

$$\begin{aligned} N(g^\sharp) = \sum_{x\in G} n(x,g) &= |g^\sharp| \sum_{\pi\in\hat{G}} \frac{\overline{\chi_\pi(g)}}{\dim \pi} \sum_{x\in G} |\chi_\pi(x)|^2 \\ &= |G||g^\sharp| \sum_{\pi\in\hat{G}} \frac{\overline{\chi_\pi(g)}}{\dim \pi} = |G||g^\sharp| \sum_{\pi\in\hat{G}} \frac{\chi_\pi(g)}{\dim \pi}, \end{aligned}$$

and it follows that

$$N(g) = |G| \sum_{\pi\in\hat{G}} \frac{\chi_\pi(g)}{\dim \pi},$$

using (4.30); this is what we wanted. □

Remark 4.4.4. (1) If we isolate the contribution of the trivial representation, we see that the number of (x,y) with $[x,y] = g$ is given by

$$|G|\Bigl(1 + \sum_{\pi\neq 1} \frac{\chi_\pi(g)}{\chi_\pi(1)}\Bigr).$$

Suppose the group G has no non-trivial one-dimensional representation (which means that the commutators *generate* G). If we apply the basic intuition of harmonic analysis, we can expect that in many circumstances the first term will dominate, and hence that many elements in G will be commutators. There are indeed many results in this direction. For instance, a recent theorem of Liebeck, O'Brien, Shalev and Tiep [**28**], confirming a striking conjecture of Ore, shows that if G is a finite non-abelian simple group, *every* element of G is a commutator. The criterion used to detect commutators in this work is the one we just proved. In Exercise 4.6.16, the reader will be invited to determine the commutators in $\mathrm{GL}_2(\mathbf{F}_p)$.

On the other hand, the reader may check (!) using software like Magma [**5**] or Gap [**14**] that there exists a perfect group of order 960, which fits into an exact sequence

$$1 \longrightarrow (\mathbf{Z}/2\mathbf{Z})^4 \longrightarrow G \longrightarrow A_5 \longrightarrow 1,$$

such that the number of actual commutators $[x,y]$ in G is not equal to $|G|$. To be more precise, this group G is isomorphic to the commutator subgroup of the group W_5 discussed below in Exercise 4.7.13, with the homomorphism to A_5 defined as the restriction to $[W_5, W_5]$ of the natural surjection $W_5 \longrightarrow \mathfrak{S}_5$. It inherits from W_5 a faithful (irreducible) representation of dimension 5 by signed permutation matrices, and it turns out that the

120 elements in the conjugacy class of the element

$$g = \begin{pmatrix} 0 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 \end{pmatrix}$$

are not commutators, as one can see that all commutators have trace in $\{-3, -2, 0, 1, 2, 5\}$. On the other hand, one can see that g *is* a commutator in W_5. (Note that, because G contains at least 481 commutators, it also follows that any $g \in G$ is the product of at most two commutators, by the following well-known, but clever, argument: in any finite group G, if S_1 and S_2 are subsets of G with $|S_i| > |G|/2$, and if $g \in G$ is arbitrary, the fact that[5] $|S_1| + |gS_2^{-1}| > |G|$ implies that $S_1 \cap gS_2^{-1} \neq \emptyset$, so that g is always of the form s_1s_2 with $s_i \in S_i$; the end of the proof of Theorem 4.7.1 in Section 4.7.1 will use an even more clever variant of this argument involving three subsets...)

(2) If we take $g = 1$ in (4.29), we see that the number of (x, y) in $G \times G$ such that $xy = yx$ is equal to

$$|G||\hat{G}| = |G||G^\sharp|.$$

The reader should attempt to prove this directly (without using characters).

The representations of a finite group G can also be used to understand other spaces of functions. We give two further examples, by showing how to construct fairly convenient orthonormal bases of functions on a quotient G/K, as well as on a given coset of a suitable subgroup.

Proposition 4.4.5 (Functions on G/K). *Let G be a finite group and let H be a subgroup of G. Let V be the space of complex-valued functions on the quotient G/H, with the inner product*

$$\langle \varphi_1, \varphi_2 \rangle_V = \frac{1}{|G/K|} \sum_{x \in G/H} \varphi_1(x)\overline{\varphi_2(x)}.$$

For $\pi : G \longrightarrow \mathrm{GL}(E)$ and $v \in E^H$, $w \in E$, define

$$\varphi_{\pi,v,w} : gH \mapsto \sqrt{\dim E}\langle \pi(g)v, w \rangle_E.$$

Then the family of functions $(\varphi_{\pi,v,w})$, where π runs over $\hat{G}$, w runs over an orthonormal basis of the space E_π of π and v runs over an orthonormal basis of the E_π^H, forms an orthonormal basis of V.

Proof. First of all, the functions $\varphi_{\pi,v,w}$ are well-defined (that is, they are functions in V), because replacing g by gh, with $h \in H$, leads to

$$\langle \pi(gh)v, w \rangle_\pi = \langle \pi(g)\pi(h)v, w \rangle_\pi = \langle \pi(g)v, w \rangle_\pi,$$

since $v \in E_\pi^H$. We observe next that if $\tilde{\varphi}_{\pi,v,w}$ denote the corresponding matrix coefficients of G, we have

$$\langle \varphi_{\pi_1,v_1,w_1}, \varphi_{\pi_2,v_2,w_2} \rangle_V = \langle \tilde{\varphi}_{\pi_1,v_1,w_1}, \tilde{\varphi}_{\pi_2,v_2,w_2} \rangle,$$

so that the family of functions indicated is, by the orthonormality of matrix coefficients, an orthonormal family in V.

[5] We use here the notation $S^{-1} = \{x^{-1} \mid x \in S\}$.

It only remains to show that these functions span V. But their total number is

$$\begin{aligned}\sum_{\pi\in\hat{G}}(\dim\pi)(\dim\pi^H)&=\sum_{\pi\in\hat{G}}(\dim\pi)\langle\operatorname{Res}_1^H\pi,\mathbf{1}_H\rangle_H\\&=\sum_{\pi\in\hat{G}}(\dim\pi)\langle\pi,\operatorname{Ind}_H^K(\mathbf{1}_H)\rangle_G\end{aligned}$$

by Frobenius reciprocity. However, for any representation ϱ of G, we have

$$\sum_{\pi\in\hat{G}}(\dim\pi)\langle\pi,\varrho\rangle_G=\sum_{\pi\in\hat{G}}(\dim\pi)n_\pi(\varrho)=\dim\varrho,$$

so that the number of functions in our orthonormal system is equal to

$$\dim\operatorname{Ind}_H^K(\mathbf{1}_H)=[G:H]=\dim(V),$$

which means that this system is in fact an orthonormal basis. □

The second case is a bit more subtle. We consider a finite group G, and a normal subgroup $H\triangleleft G$ such that the quotient A in the exact sequence

$$1\longrightarrow H\longrightarrow G\xrightarrow{\phi}A\longrightarrow 1$$

is *abelian.* Fixing $a\in A$, we want to describe an orthonormal basis of the space W of class-functions supported on the H-coset $\phi^{-1}(a)=Y\subset G$, with the inner product

$$\langle\varphi_1,\varphi_2\rangle_W=\frac{1}{|H|}\sum_{y\in Y}\varphi_1(x)\overline{\varphi_2(x)}.$$

This makes sense because H is normal in G, which implies that any coset of H is a union of conjugacy classes.

The basic starting point is that the restrictions of characters to Y still form a generating set of W (because one can extend by zero any function in W, obtaining a class function on G, which becomes a linear combination of characters). For dimension reasons, this can not be a basis (except if $H=G$). In order to extract a basis, and to attempt to make it orthonormal, we therefore need to compute the inner product in W of characters restricted to Y. In order to detect the condition $y\in Y$, we use the orthogonality of (one-dimensional) irreducible characters of the quotient group A: we have

$$\frac{1}{|A|}\sum_{\psi\in\hat{A}}\overline{\psi(\alpha)}\psi(\phi(g))=\begin{cases}0&\text{if }\phi(g)\neq\alpha,\text{ i.e., if }g\notin Y,\\1&\text{if }g\in Y,\end{cases}$$

and hence

$$\begin{aligned}\langle\chi_{\pi_1},\chi_{\pi_2}\rangle_W&=\frac{1}{|H|}\sum_{y\in Y}\chi_{\pi_1}(y)\overline{\chi_{\pi_2}(y)}\\&=\frac{1}{|H||A|}\sum_{\psi\in\hat{A}}\overline{\psi(\alpha)}\sum_{y\in Y}\psi(\phi(y))\chi_{\pi_1}(y)\overline{\chi_{\pi_2}(y)}\\&=\sum_{\psi\in\hat{A}}\overline{\psi(\alpha)}\langle(\psi\circ\phi)\otimes\pi_1,\pi_2\rangle_G\\&=\sum_{\substack{\psi\in\hat{A}\\\pi_2\simeq(\psi\circ\phi)\otimes\pi_1}}\overline{\psi(\alpha)}.\end{aligned}$$

This is more complicated than the orthogonality formula, but it remains manageable. It shows that the characters remain orthogonal on H unless we have $\pi_1 \simeq (\psi \circ \phi) \otimes \pi_2$ for some $\psi \in \hat{A}$. This is natural, because evaluating the characters, we obtain in that case

$$\chi_{\pi_1}(y) = \psi(\alpha)\chi_{\pi_2}(y)$$

for $y \in Y$, i.e., χ_{π_1} and χ_{π_2} are then *proportional.* The factor $\psi(\alpha)$ is of modulus one, and hence

$$\langle \chi_{\pi_1}, \chi_{\pi_2} \rangle_W = \langle \chi_{\pi_1}, \chi_{\pi_1} \rangle_W = \sum_{\substack{\psi \in \hat{A} \\ (\psi\circ\phi)\otimes\pi_1 \simeq \pi_1}} \overline{\psi(\alpha)}.$$

in that case. This can still be simplified a bit: if ψ occurs in the sum, we obtain

$$\psi(\alpha)\chi_{\pi_1}(y) = \chi_{\pi_1}(y)$$

for all $y \in Y$, and therefore *either* $\psi(\alpha) = 1$, or $\chi_{\pi_1}(y) = 0$ for all $y \in Y$. In the second case, the character actually vanishes on all of Y (and will not help in constructing an orthonormal basis, so we can discard it!), while in the first case, we get

$$\langle \chi_{\pi_1}, \chi_{\pi_1} \rangle_W = |\{\psi \in \hat{A} \mid (\psi \circ \phi) \otimes \pi_1 \simeq \pi_1\}|.$$

Let us denote by

$$\kappa(\pi) = |\{\psi \in \hat{A} \mid (\psi \circ \phi) \otimes \pi \simeq \pi\}| \tag{4.34}$$

the right-hand side of this formula: it is an invariant attached to any irreducible representation of G. We can then summarize as follows the discussion:

Proposition 4.4.6 (Functions on cosets). *Let G be a finite group, $H \lhd G$ a normal subgroup with abelian quotient G/H.*

For $\alpha \in A$, Y and W as defined above, an orthonormal basis of W is obtained by considering the functions

$$\varphi_\pi(y) = \frac{1}{\sqrt{\kappa(\pi)}}\chi_\pi(y)$$

for $y \in Y$, where π runs over a subset $\hat{G}_H$ defined by (1) removing from $\hat{G}$ those π such that the character of π is identically 0 on Y; (2) considering among other characters only a set of representatives for the equivalence relation

$$\pi_1 \sim_H \pi_2 \textit{ if and only if } \mathrm{Res}^G_H \pi_1 \simeq \mathrm{Res}^G_H \pi_2.$$

To completely prove this, we must simply say a few additional words to explain why the relation $\pi_1 \sim_H \pi_2$ in the statement is equivalent with the existence of $\psi \in \hat{A}$ such that $\pi_2 \simeq (\psi \circ \phi) \otimes \pi_1$: in one direction this is clear (evaluating the character, which is 1 on $H \supset \ker \psi$), and otherwise, if $\mathrm{Res}^G_H \pi_1 \simeq \mathrm{Res}^G_H \pi_2$, we apply the inner product formula with $\alpha = 0$ (so that $Y = H$) to get

$$0 \neq \langle \mathrm{Res}^G_H \pi_1, \mathrm{Res}^G_H \pi_2 \rangle_H = \sum_{\substack{\psi \in \hat{A} \\ \pi_2 \simeq (\psi\circ\phi)\otimes\pi_1}} \overline{\psi(\alpha)},$$

so that the sum can not be empty, and the existence of ψ follows. This remark means that the functions described in the statement are an orthonormal system in W. We observed at the beginning of the computation that they generate W, and hence we are done.

Exercise 4.4.7. In the situation of Proposition 4.4.6, show how to obtain an orthonormal basis of the space of all functions $Y \longrightarrow \mathbf{C}$, with respect to the inner product on $C(G)$, using restrictions of matrix coefficients. [Hint: Example 3.3.7 can be useful.]

Exercise 4.4.8. Let $\mathbf{F}_q$ be a finite field with q elements and $n \geqslant 2$ an integer.

(1) Show that taking $G = \mathrm{GL}_n(\mathbf{F}_q)$ and $H = \mathrm{SL}_n(\mathbf{F}_q)$ gives an example of the situation considered above. What is A in that case?

(2) Show that, in this case, the invariant defined in (4.34) satisfies

$$\kappa(\pi) \leqslant n$$

for any irreducible representation $\pi \in \hat{G}$.

In Exercise 4.6.5, we will give examples of groups having representations where $\kappa(\pi) \neq 1$, and also examples where the set $\hat{G}_H$ differs from $\hat{G}$ (for both possible reasons: characters vanishing on Y, or two characters being proportional on Y).

4.5. Finite abelian groups

Finite abelian groups are the easiest groups to deal with when it comes to representation theory. Since they are also very important in applications, we summarize here again the results of the previous sections, specialized to abelian groups, before discussing some features which are specific this situation.

Theorem 4.5.1 (Finite abelian groups). *Let G be a finite abelian group.*

(1) *There are exactly $|G|$ one-dimensional representations, often simply called* characters *of G, namely group homomorphisms*

$$\chi : G \longrightarrow \mathbf{C}^{\times}.$$

(2) *Let $\hat{G}$ be the set of characters of G. We have the orthogonality relations*

$$\sum_{x\in G} \chi_1(x)\overline{\chi_2(x)} = \begin{cases} |G| & \text{if } \chi_1 = \chi_2, \\ 0 & \text{if } \chi_1 \neq \chi_2, \end{cases} \tag{4.35}$$

for χ_1, $\chi_2 \in \hat{G}$,

$$\sum_{\chi\in \hat{G}} \chi(x)\overline{\chi(y)} = \begin{cases} |G| & \text{if } x = y, \\ 0 & \text{if } x \neq y, \end{cases} \tag{4.36}$$

for x, $y \in G$.

(3) *Let $\varphi : G \longrightarrow \mathbf{C}$ be any function on G. We have the Fourier decomposition*

$$\varphi = \sum_{\chi\in\hat{G}} \hat{\varphi}(\chi)\chi$$

where

$$\hat{\varphi}(\chi) = \langle \varphi, \chi\rangle = \frac{1}{|G|}\sum_{x\in G} \varphi(x)\overline{\chi(x)},$$

and the Plancherel formula

$$\sum_{\chi\in\hat{G}} |\hat{\varphi}(\chi)|^2 = \frac{1}{|G|}\sum_{x\in G} |\varphi(x)|^2.$$

The crucial feature which is specific to abelian groups is that, since all irreducible representations are of dimension 1, they form a *group* under pointwise multiplication: if χ_1, χ_2 are in $\hat{G}$, the product

$$\chi_1\chi_2 : x \mapsto \chi_1(x)\chi_2(x)$$

is again in $\hat{G}$. Similarly the inverse

$$\chi^{-1} : x \mapsto \chi(x)^{-1} = \overline{\chi(x)}$$

(where the last is because $|\chi(x)| = 1$ for all characters) is a character, and with the trivial character as neutral element, we see that $\hat{G}$ is also a group, in fact a finite abelian group of the same order as G. Its properties are summarized by:

Theorem 4.5.2 (Duality of finite abelian groups). *Let G be a finite abelian group, and $\hat{G}$ the group of characters of G, called the* dual group*.*

(1) *There is a* canonical *isomorphism*

$$\begin{cases} G & \longrightarrow & \hat{\hat{G}} \\ x & \mapsto & e_x \end{cases}$$

where e_x is the homomorphism of evaluation at x defined on $\hat{G}$, i.e.

$$e_x(\chi) = \chi(x).$$

(2) *The group $\hat{G}$ is* non-canonically *isomorphic to G.*

Proof. (1) A simple check shows that e is a group homomorphism. To show that it is injective, we must show that if $x \neq 1$, there is at least one character χ with $\chi(x) \neq 1$. This follows, for instance, from the orthogonality relation

$$\sum_{\chi} \chi(x) = 0.$$

(2) The simplest argument is to use the structure theory of finite abelian groups: there exist integers $r \geqslant 0$ and positive integers

$$d_1 \mid d_2 \mid \cdots \mid d_r$$

such that

$$G \simeq \mathbf{Z}/d_1\mathbf{Z} \times \cdots \times \mathbf{Z}/d_r\mathbf{Z}.$$

Now we observe that for a direct product $G_1 \times G_2$, there is a natural isomorphism

$$\begin{cases} \hat{G}_1 \times \hat{G}_2 \longrightarrow \widehat{G_1 \times G_2} \\ (\chi_1, \chi_2) \mapsto \chi_1 \boxtimes \chi_2, \end{cases}$$

with $(\chi_1 \boxtimes \chi_2)(x_1, x_2) = \chi_1(x_1)\chi_2(x_2)$. Indeed, this is a group homomorphism, which is quite easily seen to be injective, and the two groups have the same order.[6]

Thus we find

$$\hat{G} \simeq \widehat{\mathbf{Z}/d_1\mathbf{Z}} \times \cdots \times \widehat{\mathbf{Z}/d_r\mathbf{Z}}$$

and this means that it is enough to prove that $\hat{G} \simeq G$ when G is a finite cyclic group $\mathbf{Z}/d\mathbf{Z}$. But a homomorphism

$$\chi : \mathbf{Z}/d\mathbf{Z} \longrightarrow \mathbf{C}^\times$$

is determined uniquely by $e_1(\chi) = \chi(1)$. This complex number must be a d-th root of unity, and this means that we have an isomorphism

$$e_1 : \begin{cases} \widehat{\mathbf{Z}/d\mathbf{Z}} & \longrightarrow & \boldsymbol{\mu}_d = \{z \in \mathbf{C}^\times \mid z^d = 1\} \\ \chi & \mapsto & \chi(1), \end{cases}$$

Since the group of d-th roots of unity in $\mathbf{C}^\times$ is isomorphic (non-canonically, if $d \geqslant 3$) to $\mathbf{Z}/d\mathbf{Z}$, we are done. □

[6] It is also surjective by an application of Proposition 2.3.17.

Remark 4.5.3. In practice, one uses very often the explicit description of characters of $\mathbf{Z}/m\mathbf{Z}$ that appeared in this proof. Denoting $e(z) = e^{2i\pi z}$ for $z \in \mathbf{C}$, they are the functions of the form

$$e_a : x \mapsto e\Big(\frac{ax}{m}\Big)$$

where $x \in \mathbf{Z}/m\mathbf{Z}$ and $a \in \mathbf{Z}/m\mathbf{Z}$. In this description, of course, the exponential is to be interpreted as computed using representatives in $\mathbf{Z}$ of x and a, but the result is independent of these choices (simply because $e(k) = 1$ if $k \in \mathbf{Z}$).

Exercise 4.5.4. We have derived the basic facts about representations of finite abelian groups from the general results of this chapter. However, one can also prove them using more specific arguments. This exercise discusses one possible approach (see [**35**, VI.1]).

(1) Prove the orthogonality relation (4.35) directly.

(2) Show – without using anything else than the definition of one-dimensional characters – that if $H \subset G$ is a subgroup of a finite abelian group, and $\chi_0 \in \hat{H}$ is a character of H, there exists a character $\chi \in \hat{G}$ of G such that χ restricted to H is equal to χ_0. [Hint: Use induction on $|G/H|$.] Afterward, reprove this using Frobenius reciprocity instead, and compare with Exercise 2.3.11.

(3) Deduce from this the orthogonality relation (4.36).

(4) Deduce that $\hat{G}$ is an abelian group of the same order as G, and that the homomorphism e is an isomorphism.

Example 4.5.5. (1) Quite often, the Fourier decomposition is applied to the characteristic function $\mathbf{1}_A$ of a subset $A \subset G$. Isolating – as in the previous section – the contribution of the trivial character, we then have the expansion

$$\mathbf{1}_A(x) = \frac{|A|}{|G|} + \sum_{\substack{\chi \in \hat{G} \\ \chi \neq 1}} \alpha(\chi)\chi(x)$$

with

$$\alpha(\chi) = \sum_{x \in A} \overline{\chi(x)}.$$

The interpretation of the first term is the "probability" that an element $x \in G$, chosen uniformly at random, belongs to A. Neglecting the other terms and using only this probabilistic term is a common method to reason, on a heuristic level, about what "should" be true for certain problems.

(2) In particular, it is often interesting to use the characteristic function of the set A of d-powers in G, for some fixed d, i.e.,

$$A = \{x \in G \mid x = y^d \text{ for some } y \in G\}.$$

Since $y \mapsto y^d$ is surjective when d is coprime to the order of G, one assumes that $d \mid |G|$.

Example 4.5.6 (Dirichlet's Theorem on primes in arithmetic progressions). We sketch how Dirichlet succeeded in proving Theorem 1.2.2. Thus we have a positive integer $q \geqslant 1$ and an integer $a \geqslant 1$ coprime with q, and we want to find prime numbers $p \equiv a \,(\mathrm{mod}\, q)$.

Dirichlet's proof is motivated by an earlier argument that Euler used to reprove that there are infinitely many prime numbers, as follows: we know that

$$\lim_{\sigma \to 1} \sum_{n \geqslant 1} \frac{1}{n^\sigma} = +\infty,$$

e.g., by comparison of the series with the integral

$$\int_1^{+\infty} x^{-\sigma}dx = \frac{1}{\sigma-1}, \qquad \text{for } \sigma > 1.$$

On the other hand, exploiting the unique factorization of positive integers in products of primes, Euler showed that

$$\sum_{n\geqslant 1}\frac{1}{n^\sigma} = \prod_p (1-p^{-\sigma})^{-1} \tag{4.37}$$

for $\sigma > 1$, where the infinite product is over all prime numbers. (It is defined as the limit, as $x \to +\infty$, of the partial products

$$\prod_{p\leqslant x}(1-p^{-\sigma})^{-1} = \prod_{p\leqslant x}\sum_{k\geqslant 0} p^{-k\sigma} = \sum_{n\in P(x)} n^{-\sigma}$$

where $P(x)$ is the sum of all positive integers with no prime divisor $> x$, and the unique factorization of integers has been used in the last step; thus the absolute convergence of the series on the left is then enough to justify the equality (4.37).

Now obviously, if there were only finitely many primes, the right-hand side of the formula would converge to some fixed real number as $\sigma \to 1$, which contradicts what we said about the left-hand side. Hence there are infinitely many primes.

An equivalent way to conclude is to take the logarithm on both sides; denoting

$$\zeta(\sigma) = \sum_{n\geqslant 1} n^{-\sigma}$$

for $\sigma > 1$, we have

$$\log\zeta(\sigma) \to +\infty$$

on the one-hand, and on the other hand

$$\log\zeta(\sigma) = -\sum_p \log(1-p^{-\sigma}) = \sum_p p^{-\sigma} + \sum_{p,k\geqslant 2} k^{-1}p^{-k\sigma} = \sum_p p^{-\sigma} + O(1)$$

as $\sigma \to 1$ (where we have used the power series expansion

$$\log\Bigl(\frac{1}{1-x}\Bigr) = \sum_{k\geqslant 1}\frac{x^k}{k}$$

which converges absolutely and uniformly on compact sets for $|x| < 1$.) Thus Euler's argument can be phrased as

$$\lim_{\sigma\to 1}\sum_p p^{-\sigma} = +\infty.$$

Using this, it is rather tempting (isn't it?) to try to analyze either the product or the sum

$$\prod_{p\equiv a \,(\mathrm{mod}\, q)} (1-p^{-\sigma})^{-1}, \qquad \sum_{p\equiv a\,(\mathrm{mod}\, q)} p^{-\sigma}$$

similarly, and to show that, as $\sigma \to 1$, these functions tend to $+\infty$ (as before, they differ at most by a bounded function). But if we expand the product, we do *not* get the "obvious" series

$$\sum_{\substack{n\geqslant 1\\ n\equiv a\,(\mathrm{mod}\, q)}} n^{-\sigma},$$

because there is no reason that the primes dividing an integer congruent to a modulo q should have the same property (also, if a is not $\equiv 1 \pmod{q}$, the product of primes congruent to a modulo q is not necessarily $\equiv a \pmod{q}$): e.g., $35 = 7 \times 5$ is congruent to 3 modulo 4, but $5 \equiv 1 \pmod{4}$.

In other words, we are seeing the effect of the fact that the characteristic function of $a \in \mathbf{Z}/q\mathbf{Z}$, which is used to select the primes in the product or the series, is *not* multiplicative. Dirichlet's idea is to use, instead, some functions on $\mathbf{Z}/q\mathbf{Z}$ which *are* multiplicative, and to use them to recover the desired characteristic function.

A *Dirichlet character* modulo q is then defined to be a map

$$\chi \,:\, \mathbf{Z} \longrightarrow \mathbf{C}$$

such that $\chi(n) = 0$ if n is not coprime to q, and otherwise $\chi(n) = \chi_*(n \pmod{q})$ for some character of the multiplicative group of invertible residue classes modulo q:

$$\chi_* \,:\, (\mathbf{Z}/q\mathbf{Z})^\times \longrightarrow \mathbf{C}^\times.$$

It follows that $\chi(nm) = \chi(n)\chi(m)$ for all n, $m \geqslant 1$ (either because both sides are 0 or because χ_* is a homomorphism), and from the orthogonality relation we obtain

$$\sum_{\chi \pmod{q}} \overline{\chi(a)}\chi(n) = \begin{cases} |(\mathbf{Z}/q\mathbf{Z})^\times| = \varphi(q), & \text{if } n \equiv a \pmod{q} \\ 0, & \text{otherwise,} \end{cases}$$

for $n \geqslant 1$ (because a is assumed to be coprime to q), the sum ranging over all Dirichlet characters modulo q, which correspond exactly to the characters of the group $(\mathbf{Z}/q\mathbf{Z})^\times$.

This is the first crucial point: the use of "suitable harmonics" to analyze the characteristic function of a residue class. Using it by summing over n, we obtain the identity

$$\sum_{p \equiv a \pmod{q}} p^{-\sigma} = \frac{1}{\varphi(q)} \sum_{\chi \pmod{q}} \overline{\chi(a)} \sum_p \chi(p) p^{-\sigma},$$

while, for each χ, the multiplicativity leads to an analogue of the Euler product:

$$\sum_{n \geqslant 1} \chi(n) n^{-\sigma} = \prod_p (1 - \chi(p) p^{-\sigma})^{-1}.$$

As now classical, we denote by $L(\sigma, \chi)$ the function in this last formula. By the same reasoning used for $\zeta(\sigma)$, it satisfies

$$\log L(\sigma, \chi) = \sum_p \chi(p) p^{-\sigma} + O(1)$$

as $\sigma \to 1$. We have therefore

$$\sum_{p \equiv a \pmod{q}} p^{-\sigma} = \frac{1}{\varphi(q)} \sum_{\chi \pmod{q}} \overline{\chi(a)} \log L(\sigma, \chi) + O(1),$$

for all $\sigma > 1$, and now the idea is to imitate Euler by letting $\sigma \to 1$ and seeing a divergence emerge on the right-hand side, which then implies that the series on the left can not have only finitely many non-zero terms.

On the right-hand side, for the character χ_0 corresponding to $\chi_* = 1$, we have

$$L(\sigma, \chi_0) = \prod_{p \nmid q} (1 - p^{-\sigma})^{-1}$$

(the primes dividing q have $\chi_0(p) = 0$), which therefore satisfies

$$\log L(\sigma, \chi_0) = \log(1/(\sigma - 1)) + O(1)$$

as $\sigma \to 1$ – only the finitely many terms at $p \mid q$ make this different from Euler's case of $\zeta(\sigma)$. This contribution therefore diverges, and we see that Dirichlet's Theorem follows from the second crucial ingredient: the fact that for a Dirichlet character χ associated to a character $\chi_* \neq 1$, the function

$$L(\sigma, \chi)$$

converges to a *non-zero* value as $\sigma \to 1$, so that its logarithm also has a limit. Showing that the function converges to some complex number is not too difficult; however, proving that this complex number is non-zero is much more subtle. Since this has little to do with representation theory, we refer, e.g., to [**35**, Ch. 6] for a very careful presentation of the details.

Exercise 4.5.7 (Burnside's inequality for cyclic groups). Although, much of the time, one deals with irreducible characters of finite abelian groups, higher-dimensional representations do sometimes occur. Here is one result which is used in the proof of the result concerning irreducible representations of degree at least 2 of (necessarily) non-abelian group (see Proposition 4.7.11 below).

For a finite cyclic group $G = \mathbf{Z}/m\mathbf{Z}$ with $m \geqslant 1$, we let $G^* \subset G$ be the set of generators of G. The goal is to prove that if ϱ is any finite-dimensional representation of G, we have

$$\sum_{x \in G^*} |\chi_\varrho(x)|^2 \geqslant |G^*| \tag{4.38}$$

unless $\chi_\varrho(x) = 0$ for all $x \in G^*$.

(1) If you know Galois theory, prove this directly. [Hint: Use the arithmetic-geometric mean inequality.]

The next steps present an alternative argument which does not require Galois theory.

(2) Show that there exists a non-negative quadratic form Q_m in m variables, denoted $\boldsymbol{n} = (n(a))_{a \in \mathbf{Z}/m\mathbf{Z}}$, such that

$$\sum_{x \in G^*} |\chi_\varrho(x)|^2 = Q_m(\boldsymbol{n})$$

for any representation ϱ, where the coordinate $n(a)$ of $\boldsymbol{n} = (n(a))$ is the multiplicity of the irreducible character

$$x \mapsto e\Bigl(\frac{ax}{m}\Bigr)$$

of $\mathbf{Z}/m\mathbf{Z}$.

(3) Show that if

$$m = \prod_{p \mid m} p^{k_p}$$

is the prime factorization of m, we have

$$Q_m = \bigotimes_{p \mid m} Q_{p^{k_p}}.$$

(4) Show that for p prime and for any quadratic form Q' of rank $d \geqslant 1$, we have

$$(Q_p \otimes Q')(\boldsymbol{n}) = \frac{1}{2} \sum_{a,b \in \mathbf{Z}/p\mathbf{Z}} Q'(n(a) - n(b))$$

for any $\boldsymbol{n} = (n(a)) \in (\mathbf{Z}^d)^p$. [Hint: It may be useful to start with $Q'(n) = n^2$ of rank 1.]

(5) For Q' as above, non-negative, let $s(Q')$ denote the smallest non-zero value of $Q'(\boldsymbol{n})$ for $\boldsymbol{n} \in \mathbf{Z}^d$. Show that for any quadratic form Q' of rank $d \geqslant 1$, we have

$$s(Q_p \otimes Q') = (p-1)s(Q').$$

(6) For $k \geqslant 2$ and p prime, show that there exists a quadratic form Q' of rank p^{k-1} such that $Q_{p^k} = Q_p \otimes Q'$ and $s(Q') = p^{k-1}$. Then prove (4.38).

4.6. The character table

The "character table" of a finite group G is the name given to the matrix $(\chi_\varrho(c))_{\varrho,c}$ which gives the values of all the (complex) irreducible characters χ_ϱ of G at all conjugacy classes $c \in G^\sharp$. In particular, it is a square matrix which determines the irreducible characters as class functions, and hence encapsulates in theory all the information given by representation theory for the group G, over $\mathbf{C}$ at least. It is typically represented as a square matrix with rows given by the irreducible characters (in some order) and columns indexed by the conjugacy classes.

Example 4.6.1. A very simple example is the character table of $G = \mathfrak{S}_3$:

	1	(12)	(123)
1	1	1	1
ε	1	-1	1
ϱ_2	2	0	-1

Table 4.1. Character table of $\mathfrak{S}_3$

Here the top line, as well as the leftmost row, simply recall the chosen ordering of the conjugacy classes and characters. For the former, this is usually fairly self-explanatory, but for the characters, one often wants – if possible – a description of an actual representation which has the character values given in the row (if only to check that it is correctly described!)

This might be a complicated matter, but for this example, this is simple: **1** is the trivial representation, $\varepsilon : \mathfrak{S}_3 \longrightarrow \mathbf{C}^\times$ is the signature, and ϱ_2 is the 2-dimensional representation acting on

$$E = \{(x,y,z) \in \mathbf{C}^3 \mid x+y+z=0\}.$$

Indeed, it is not hard to check that the character values are correct, and one can check that ϱ_2 is irreducible by computing the norm $\langle \chi_{\varrho_2}, \chi_{\varrho_2} \rangle = 1$. For the latter, note that one needs to know the order of the conjugacy classes to weigh properly the indicated values (see Remark 4.4.2). This extra information is often indicated in parallel with the character table, but it can in fact be recovered[7] from it using (4.28): if we fix a class $c \in G^\sharp$ and take $h \in c$ there, we see that

$$|c| = \frac{|G|}{\sum_\varrho |\chi_\varrho(h)|^2}. \tag{4.39}$$

For instance, taking $c = (12)$ and $c' = (123)$ for $\mathfrak{S}_3$, we get

$$|c| = \frac{6}{1^2+1^2} = 3, \qquad |c'| = \frac{6}{1^2+1^2+1^2} = 2.$$

[7] This, of course, assumes the full character table is indeed known...

as it should.

Before we discuss which information concerning a group can be extracted from the character table, it is interesting to ask: can we always compute it? To be (a bit) more precise:

Proposition 4.6.2 (The character table is computable). *Let G be a finite group, given in such a way that: one can enumerate all elements of G and one can compute the group law and the inverse.*[8] *Then there is an algorithm that will terminate after some time by listing the character table of G.*

This is a fairly poor version of computability: we make no claim, or guarantee, about the amount of time the algorithm will require (an estimate can be obtained from the argument, but it will be very bad).

Proof. First of all, by enumerating all pairs of elements and computing all xyx^{-1}, one can make the list of all the conjugacy classes of G, and find means to associate its conjugacy class to any element of G.

The idea is then to see that the regular representation can be decomposed into isotypic components. Indeed, first of all, the regular representation is computable: a basis of $C(G)$ is given by characteristic functions of a single point, and the action on the basis vectors is computable from the inverse map. Moreover, decomposing an arbitrary $f \in C(G)$ in this basis is immediate.

It is then enough to find an algorithm to compute the decomposition

$$C(G) = \bigoplus_{\pi \in \hat{G}} M(\pi)$$

of the regular representation into isotypic components, in the sense of giving a list of bases of the spaces $M(\pi)$. Indeed, given a subspace M among these, one can then compute the corresponding character by

$$\chi(g) = \frac{1}{\dim(M)} \operatorname{Tr}(\operatorname{reg}(g)|_M).$$

Now the crucial step: the subspaces $M(\pi)$ are characterized as the common eigenspaces of all operators $\operatorname{reg}(a)$ where $a \in Z(k(G))$ is an element in the center of the group algebra, or equivalently as the common eigenspaces of the operators $\operatorname{reg}(a_c)$, where c runs over the conjugacy classes of G and

$$a_c = \sum_{g \in c} g$$

as before. In other words, assume a subspace $M \subset C(G)$ has the property that there exist eigenvalues $\lambda_c \in \mathbf{C}$, defined for all c, such that

$$M = \{v \in C(G) \mid \operatorname{reg}(a_c)v = \lambda_c v \text{ for all } c \in G^\sharp\} \ ; \tag{4.40}$$

then M is one of the $M(\pi)$, and conversely.

If this is granted, we proceed as follows: list the conjugacy classes, and for each of them, find the eigenvalues and eigenspaces of the operator $\operatorname{reg}(a_c)$ (by finding bases for them, using linear algebra, which is eminently computable), then compute their intersections, and list the resulting subspaces: they are the isotypic components $M(\pi)$.

[8] As an example of the subtleties that may be involved, note that having a generating set is not enough: one must be able to say whether two arbitrary products of elements from such a set are equal or not.

We now check the claim. Let M be a common eigenspace of the $\mathrm{reg}(a_c)$ given by (4.40). Writing irreducible characters[9] as combinations of the a_c's, and taking linear combinations, it follows that M is also contained in an eigenspace of the isotypic projectors e_π, for any π. But these are $M(\pi)$, for the eigenvalue 1, and the sum of the other isotypic components, for the eigenvalue 0. The sum of the e_π is the identity, so there exists some $\pi \in \hat{G}$ such that the eigenvalue of e_π on M is 1. This means that $M \subset M(\pi)$. But $M(\pi)$ itself is contained in a common eigenspace:

$$M(\pi) \subset \Big\{ v \in C(G) \mid \mathrm{reg}(a_c)v = \frac{|c|\chi_\pi(c)}{|G|} v \text{ for all } c \Big\},$$

by Proposition 4.7.11. This shows that the λ_c must coincide with $\frac{|c|\chi_\pi(c)}{|G|}$. But then

$$M \subset M(\pi) \subset M,$$

and these inclusions must be equalities! □

This algorithm is not at all practical if G is large, but at least it shows that, by trying to get information *from* the character table, we can not be building castles completely in the air!

Note that besides this fact, the proof has led to the characterization

$$M(\pi) = \Big\{ v \in C(G) \mid \mathrm{reg}(a)v = \omega_\pi(a)v \text{ for all } a \in Z(k(G)) \Big\}$$

of the isotypic components of the regular representation, which is of independent interest.

Exercise 4.6.3. Proposition 4.6.2 shows how to compute, in principle, the character table of a finite group G. Explain how one can also, in principle, write matrix representations $\pi^{\boldsymbol{m}} : G \longrightarrow \mathrm{GL}_{\dim(\pi)}(\mathbf{C})$ for each irreducible representation $\pi \in \hat{G}$.

4.6.1. Some features of the character table. We present here some of the information that can be derived from the character table of a group, if it is known (other examples are given for instance in [**19**, Ch. 2]). Like in Proposition 4.6.2, we do not attempt to measure the actual computational efficiency of the procedures we describe, many of which are quite impractical when implemented directly. In the next sections, we will compute the character tables of some concrete groups in detail.

– As already noticed, the sizes of the conjugacy classes, or equivalently the sizes of the centralizers of $C_G(g) = |G|/|g^\sharp|$ of elements in G, can be computed from the character table using the formula (4.39).

– The kernel of an irreducible representation ϱ can be determined from its row of the character table, because of the following:

Proposition 4.6.4 (Size of the character values). *Let G be an arbitrary group and ϱ a finite-dimensional unitary, or unitarizable, representation of G. For $g \in G$, we have*

$$|\chi_\varrho(g)| \leqslant \dim \varrho, \tag{4.41}$$

and there is equality if and only if $\varrho(g)$ is a scalar. *In particular, $\varrho(g) = 1$ if and only if $\chi_\varrho(g) = \chi_\varrho(1) = \dim(\varrho)$.*

[9] We are allowed now to use characters as theoretical tools to check that the algorithm works!

Proof. Since ϱ is unitary, the eigenvalues of $\varrho(g)$ are of modulus 1, and hence its trace, the value $\chi_\varrho(g)$, is at most $\dim(\varrho)$. Moreover, by the equality case of the triangle inequality, there can be equality only if all eigenvalues are equal, which means (since $\varrho(g)$ is diagonalizable) that $\varrho(g)$ is the multiplication by this common eigenvalue.

Finally, in that case, we can compute the eigenvalue as $\chi_\varrho(g)/\dim(\varrho)$, and hence the latter is equal to 1 if and only if $\varrho(g)$ is the identity. □

Note, however, that in general a character, even for a faithful representation, has no reason to be *injective* (on G or on conjugacy classes): the regular representation gives an example of this (it is faithful but, for $|G| \geqslant 3$, its character is not an injective function on G (Example 2.7.34)). Another type of "failure of injectivity" related to characters is described in Exercise 4.6.14.

– More generally, all normal subgroups of G can be computed using the character table, as well as their possible inclusion relations. Indeed, first one can find the kernels of the irreducible representations using the lemma and the character table. Then, if $N \triangleleft G$ is any normal subgroup, we have

$$N = \ker \varrho_N$$

where ϱ_N is the permutation representation associated to the left-action of G on G/N (indeed, if e_{xN} are the basis vectors for the space of ϱ_N, to have $g \in \ker \varrho_N$ means that $ge_{xN} = e_{xN}$ for all $x \in G$, i.e., $gxN = xN$ for all x, so $g \in N$, and the converse follows because N is normal) and if we denote by

$$X = \{\varrho \in \hat{G} \mid \langle \varrho, \mathrm{Ind}_N^G(\mathbf{1})\rangle \not\geqslant 1\}$$

the set of those irreducible representations which occur in the induced representation, it follows that

$$N = \bigcap_{\varrho \in X} \ker \varrho.$$

Thus, to determine all normal subgroups of G, from the character table, one can first list all kernels of irreducible representations, and then compute all intersections of finitely many such subgroups. In particular, once the conjugacy classes which form a normal subgroup are known, its order can of course be computed by summing their size. Note that, on the other and, there is no way to determine *all* subgroups of G from the character table (see Exercise 4.6.6).

– The character table of a quotient G/N of G by a normal subgroup $N \triangleleft G$ can also be determined from the character table, since irreducible representations of G/N correspond bijectively to those irreducible representations of G where $N \subset \ker \varrho$.

– Whether G is solvable can then, in principle, be determined from the character table. Indeed, we can determine whether G contains a proper normal subgroup N with abelian quotient G/N (one can check if a group is abelian by checking that all irreducible representations have dimension 1); if $N = G$, then G is abelian, hence solvable, and if N does not exist, the group is not solvable. Otherwise, we can iterate with G replaced by G/N.

– It is also natural to ask what *can not* be determined from the character table. At first, one might hope that the answer would be "nothing at all!", i.e., that it may be used to characterize the group up to isomorphism. This is not the case, however,[10] as we will show with a very classical example of two non-isomorphic groups with distinct character

[10] The reader should *not* add "unfortunately": there are no unfortunate events in mathematics...

tables, in Exercise 4.6.6 below. It will follow, in particular, that the character table can not be used to determine *all* subgroups of a finite group.

4.6.2. A nilpotent group. In order of structural complexity, after abelian groups come (non-abelian) nilpotent groups. We recall (see, e.g., [**33**, Ch. 5]) that G is nilpotent if, for some k, we have $G_k = 1$, where the sequence of subgroups G_k is defined inductively by

$$G_1 = G, \qquad G_{k+1} = [G_k, G].$$

A good example is given by the family of finite Heisenberg groups H_p defined by

$$H_p = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \mid x, y, z \in \mathbf{F}_p = \mathbf{Z}/p\mathbf{Z} \right\}$$

for p prime. Indeed, this is a p-group since $|H_p| = p^3$, and any finite p-group is nilpotent (see, e.g., [**33**, Th. 5.33]).

We will construct the character table of the groups H_p. We first gain some insight in the structure of the group H_p by computing its conjugacy classes. First of all, we will denote by the shorthand

$$\{x, y, z\}_H = \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix}$$

the elements of H_p. Then straightforward computations yield the product formula

$$\{x, y, z\}_H \{a, b, c\}_H = \{x + a, y + b, xb + z + c\}_H$$

the conjugacy formula

$$\{x, y, z\}_H \{a, b, c\}_H \{x, y, z\}_H^{-1} = \{a, b, xb - ya + c\}_H,$$

as well as the commutator relation

$$[\{x, y, z\}_H, \{a, b, c\}_H] = \{0, 0, xb - ya\}_H.$$

The last formula shows that

$$[H_p, H_p] = Z = \{\{0, 0, z\}_H \mid z \in \mathbf{F}_p\},$$

is the *center* of H_p. Each element of Z is a one-element conjugacy class in H_p; on the other hand, if $(a, b) \neq (0, 0)$, the conjugacy formula shows that, for any fixed $c \in \mathbf{F}_p$, the conjugacy class of $\{a, b, c\}_H$ is simply

$$X_{a,b} = \{\{a, b, z\}_H \mid z \in \mathbf{F}_p\}$$

(because the image of $(x, y) \mapsto xb - ya + c$ is all of $\mathbf{F}_p$ in that case, as the image of a non-constant affine map.) We have found all conjugacy classes now:

- There are p central conjugacy classes of size 1;
- There are $p^2 - 1$ conjugacy classes $X_{a,b}$ of size p.

In particular, the character table of H_p has p^2+p-1 rows and columns. To start filling it up, a good first step is to determine all one-dimensional representations χ. Not only is it a beginning to the table, but also one can then later produce new representations by considering the products $\varrho \otimes \chi$ of a "brand new" representation ϱ with the one-dimensional ones.

The one-dimensional representations are determined by computing the derived group $H_p/[H_p, H_p] = H_p/Z$, and we see here that we have an isomorphism

$$\begin{cases} H_p/Z \longrightarrow \mathbf{F}_p^2 \\ \{a, b, c\}_H \mapsto (a, b). \end{cases}$$

Thus we have p^2 distinct one-dimensional representations of H_p given by

$$\chi_{\psi_1,\psi_2} : \{a, b, c\}_H \mapsto \psi_1(a)\psi_2(b)$$

where ψ_1, ψ_2 are two characters of $\mathbf{F}_p$.

This now leaves us to find $p^2 + p - 1 - p^2 = p - 1$ irreducible representations, about which we know that the sum of the squares of their dimensions must be

$$|G| - p^2 = p^2(p - 1).$$

By comparison, it is very tempting to think that each of those new representations should be of dimension p. (Indeed, if we also use the fact that their dimension divides $|H_p| = p^3$, and hence must be a power of p, and not 1, this is the only possibility, as a representation of dimension p^2 would already have $(\dim \varrho)^2 = p^4 > p^2(p - 1)$...)

One of the most common ways of finding irreducible representations is to try to use induction to construct them, or at least to construct representations which contain "new" irreducibles. In particular, inducing one-dimensional representations of a subgroup can be quite efficient. In the case of H_p, if we want to find representations of dimension p, we can look for a subgroup of index p; for instance, we consider

$$K = \{\{0, y, z\}_H \mid y, z \in \mathbf{F}_p\} = \left\{ \begin{pmatrix} 1 & 0 & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \right\} \subset H_p.$$

We see that $K \simeq \mathbf{F}_p^2$; thus we fix a one-dimensional character ψ of K given by characters ψ_1, ψ_2 of $\mathbf{F}_p$ such that

$$\psi(\{0, y, z\}_H) = \psi_1(y)\psi_2(z),$$

and consider

$$\varrho = \mathrm{Ind}_K^{H_p}(\psi).$$

We now proceed to compute the character of this representation, using the formula (2.45). We need for this a set T of representatives of H_p/K, and we can take

$$T = \{t(x) \mid t(x) = \{x, 0, 0\}_H \text{ with } x \in \mathbf{F}_p\}.$$

Then the character of ϱ is given by

$$\chi_\varrho(g) = \sum_{\substack{x\in\mathbf{F}_p \\ t(x)gt(x)^{-1}\in K}} \psi(t(x)gt(x)^{-1}).$$

But for $g = \{a, b, c\}_H$, we have

$$t(x)gt(x)^{-1} = \begin{pmatrix} 1 & a & xb + c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix}$$

so that we see already that $\chi_\varrho(g) = 0$ if $a \neq 0$, i.e., if $g \notin K$. If $a = 0$, on the other hand, the condition $t(x)gt(x)^{-1}$ is always satisfied, and thus

$$\chi_\varrho(g) = \sum_{x\in\mathbf{F}_p} \psi_1(b)\psi_2(xb + c) = \psi_1(b) \sum_{x\in\mathbf{F}_p} \psi_2(xb + c).$$

If $b \neq 0$, sending x to $xb + c$ is a bijection of $\mathbf{F}_p$, and the result is therefore

$$\chi_\varrho(g) = \begin{cases} p\psi_1(b) & \text{if } \psi_2 = \mathbf{1} \\ 0 & \text{if } \psi_2 \neq \mathbf{1}, \end{cases}$$

while for $b = 0$, which means $g = \{0, 0, c\}_H \in Z$, we have

$$\chi_\varrho(g) = p\psi_2(c).$$

Hence there are two cases for χ_ϱ, depending on whether ψ_2 is trivial or not:

	$\{0, 0, c\}_H$	$\{0, b, \star\}_H$, $b \neq 0$	$\{a, b, \star\}_H$, $a \neq 0$
$\psi_2 = \mathbf{1}$	p	$p\psi_1(b)$	0
$\psi_2 \neq \mathbf{1}$	$p\psi_2(c)$	0	0

The middle column concerns $p-1$ non-central conjugacy classes, and the last concerns the remaining $p^2 - p$, each having p elements. Thus the respective squared norms in the two cases are

$$\frac{1}{p^3}\Bigl(p \times p^2 + p^2 \times (p-1) \times p\Bigr) = p,$$

when $\psi_2 = \mathbf{1}$ and

$$\frac{p \times p^2}{p^3} = 1$$

when $\psi_2 \neq \mathbf{1}$. Thus, whenever $\psi_2 \neq \mathbf{1}$, we have an irreducible representation. Moreover, the character values in that case show that ϱ is then independent of the choice of ψ_1, up to isomorphism, while looking at the values at the center, we see that inducing using different choices of ψ_2 leads to different representations of H_p. In other words, the $p-1$ representations

$$\varrho_{\psi_2} = \operatorname{Ind}_K^{H_p}(\psi), \qquad \psi(b, c) = \psi_2(c),$$

with ψ_2 non-trivial, give the remaining $p-1$ irreducible representations of H_p of dimension p.

We can then present the full character table as follows, where the sole restriction is that ψ_2 in the last row should be non-trivial:

	$\{0, 0, c\}_H$	$\{a, b, \star\}_H$, $(a, b) \neq (0, 0)$
χ_{ψ_1,ψ_2}	1	$\psi_1(a)\psi_2(b)$
ϱ_{ψ_2}	$p\psi_2(c)$	0

Table 4.2. Character table of H_p

4.6.3. Some solvable groups. Pursuing towards greater group-theoretic complexity, it is natural to consider some non-nilpotent solvable groups. Here a good example to handle is the family

$$B_p = \left\{ \begin{pmatrix} x & t \\ 0 & y \end{pmatrix} \mid t \in \mathbf{F}_p,\ x, y \in \mathbf{F}_p^\times \right\}$$

where p is a prime number. Thus $|B_p| = p(p-1)^2$, and the group is solvable because we have a surjective homomorphism

$$\begin{cases} B_p \longrightarrow (\mathbf{F}_p^\times)^2 \\ \begin{pmatrix} x & t \\ 0 & y \end{pmatrix} \mapsto (x, y), \end{cases} \tag{4.42}$$

with abelian kernel

$$U = \left\{ \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \mid t \in \mathbf{F}_p \right\} \simeq \mathbf{F}_p,$$

i.e., B_p is an extension of abelian groups.

As before, we compute the conjugacy classes, using the formula

$$\begin{pmatrix} x & t \\ 0 & y \end{pmatrix} \begin{pmatrix} a & u \\ 0 & b \end{pmatrix} \begin{pmatrix} x & t \\ 0 & y \end{pmatrix}^{-1} = \begin{pmatrix} a & y^{-1}\{t(b-a)+xu\} \\ 0 & b \end{pmatrix}.$$

We consider the middle matrix to be fixed, and we look for its conjugacy class. If $b \neq a$ the top-left coefficient can take any value when varying x, y and t, in fact even with $x = y = 1$, and this gives us $(p-1)(p-2)$ conjugacy classes of size p. If $a = b$, there are two cases: (1) if $u \neq 0$, we can get all *non-zero* coefficients, thus we have $p-1$ conjugacy classes of size $p-1$; (2) if $u = 0$, then in fact the matrix is scalar and its conjugacy class is a single element.

To summarize, there are:

- $p-1$ central conjugacy classes of size 1;
- $p-1$ conjugacy classes with representatives
$$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$$
of size $p-1$;
- $(p-1)(p-2)$ conjugacy classes with representatives
$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix},$$
of size p.

The number of conjugacy classes, and hence of irreducible representations, is now

$$(p-1)(p-2) + 2(p-1) = p(p-1).$$

We proceed to a thorough search... First, as in the previous section, we can easily find the one-dimensional characters; the commutator formula

$$\begin{pmatrix} x & t \\ 0 & y \end{pmatrix} \begin{pmatrix} a & u \\ 0 & b \end{pmatrix} \begin{pmatrix} x & t \\ 0 & y \end{pmatrix}^{-1} \begin{pmatrix} a & u \\ 0 & b \end{pmatrix}^{-1} = \begin{pmatrix} 1 & \text{(something)} \\ 0 & 1 \end{pmatrix}$$

shows that the morphism (4.42) factors in fact through an isomorphism

$$B_p/[B_p, B_p] \simeq (\mathbf{F}_p^\times)^2.$$

Thus we have $(p-1)^2$ one-dimensional representations

$$\varrho(\chi_1, \chi_2) : \begin{pmatrix} a & u \\ 0 & b \end{pmatrix} \mapsto \chi_1(a)\chi_2(b) \tag{4.43}$$

where χ_1 and χ_2 are one-dimensional characters of $\mathbf{F}_p^\times$. Subtracting, we see that we require

$$p(p-1) - (p-1)^2 = p-1$$

other irreducible representations, and that the sums of the squares of their dimensions must be

$$|B_p| - (p-1)^2 = p(p-1)^2 - (p-1)^2 = (p-1)^3.$$

This time, the natural guess is that there should be $p-1$ irreducible representations, each of dimension $p-1$, as this would fit the data very well... (Note also that $p-1 \mid |B_p|$, as we know it should).

This time, we will find these representations using a slightly different technique than induction. Namely, we consider the representation attached to a natural permutation representation of B_p: let X_p be the set of all lines (passing through the origin) in $\mathbf{F}_p^2$, on which B_p acts naturally (an element g acts on a line by mapping it to its image under the associated linear map $\mathbf{F}_p^2 \longrightarrow \mathbf{F}_p^2$). By definition, the line $\mathbf{F}_p e_1$ spanned by $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ is fixed by all elements in B_p, and thus we can consider the permutation representation associated to the action on $Y_p = X_p - \{\mathbf{F}_p e_1\}$. This set has order p (it contains the "vertical" line with equation $x = 0$ and the lines $y = \lambda x$ where $\lambda \in \mathbf{F}_P^\times$), and thus the associated permutation representation π has dimension p. This is not the right dimension, but we know that a permutation representation of this type always contains the trivial representation, represented by the invariant element which is the sum of the basis vectors. Thus we have a representation τ of dimension $p-1$ on the space

$$E = \{(x_\ell)_{\ell \in Y_p} \mid \sum_\ell x_\ell = 0\} \subset \mathbf{C}^{Y_p}.$$

We proceed to compute its character. This is easy because

$$\chi_\tau = \chi_\pi - 1,$$

and we know that for a permutation representation, such as π, we have

$$\chi_\pi(g) = |\{\ell \in Y_p \mid g \cdot \ell = \ell\}|,$$

the number of fixed points of the permutation associated to an element, which we can compute simply by looking at the conjugacy classes described above:

- If g is central, it fixes every line, so $\chi_\pi(g) = p$, and $\chi_\tau(g) = p - 1$;
- If
$$g = \begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix},$$
there is no fixed point, since this would correspond to an eigenvector of this matrix independent from e_1, whereas the matrix is not diagonalizable. Hence $\chi_\pi(g) = 0$, and $\chi_\tau(g) = -1$;
- If
$$g = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix},$$
with $a \neq b$, there is a unique fixed point (here, the line spanned by the second basis vector e_2); thus $\chi_\pi(g) = 1$ and $\chi_\tau(g) = 0$.

We summarize the character of τ:

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$, $a \neq b$
τ	$p-1$	-1	0

What is the squared norm of this character? We find

$$\langle \chi_\tau, \chi_\tau \rangle = \frac{1}{p(p-1)^2}\Big((p-1)^2 \times (p-1) + (p-1) \times (p-1)\Big) = 1$$

so that it *is* indeed irreducible, of dimension $p - 1$. This is just one representation, but we can "twist" using one-dimensional characters: for $\chi = \varrho(\chi_1, \chi_2)$ as in (4.43), we find that the character values of $\tau \otimes \chi$ are:

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$, $a \neq b$
$\tau \otimes \chi$	$(p-1)\chi_1(a)\chi_2(a)$	$-\chi_1(a)\chi_2(a)$	0

These are all irreducible representations, but they depend only on the product character $\chi_1\chi_2$ of $\mathbf{F}_p^\times$, and thus there are only $p-1$ different irreducible representations of dimension $p-1$ that arise in this manner.

We have now found the right number of representations. We summarize all this in the character table, using the characters $\varrho(\chi, \mathbf{1})$ to obtain the $(p-1)$-dimensional representations:

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$, $a \neq b$
$\varrho(\chi_1, \chi_2)$	$\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(b)$
$\tau \otimes \varrho(\chi, \mathbf{1})$	$(p-1)\chi(a)$	$-\chi(a)$	0

Table 4.3. Character table of B_p

Exercise 4.6.5 (Dihedral groups). The *dihedral* groups D_n, of order $2n$, form another well-known family of solvable groups; these can be defined either as the subgroup of isometries of $\mathbf{R}^2$ fixing (globally, not pointwise) a regular n-sided polygon centered at the origin, or as the group generated by a normal cyclic subgroup $C_n \simeq \mathbf{Z}/n\mathbf{Z}$ of order 2 and an element $i \in D_n - C_n$ of order 2 such that

$$ixi^{-1} = ixi = x^{-1}$$

for $x \in C_n$ (geometrically, C_n corresponds to rotations, generated by the rotation of angle $2\pi/n$, and i to an orientation-reversing isometry.)

(1) Find the character table for D_n. [Hint: They are slightly different when n is even or odd; see, e.g., [**34**, §5.2] for the details.]

(2) In the notation of Exercise 4.3.21, show that if $p \geqslant 0$ is such that

$$A_p(C_n) \leqslant A_p(D_n),$$

for n large enough, then necessarily $p \geqslant 1$.

(3) Show that if n is odd, there exist distinct irreducible representations of D_n which are proportional on the non-trivial coset $Y = D_n - C_n$ of C_n in D_n, that there exist irreducible representations with character identically zero on Y, and with the invariant

$$\kappa(\pi) = |\{\psi \in \hat{A} \mid (\psi \circ \phi) \otimes \pi \simeq \pi\}|$$

not equal to 1, where $\phi : D_n \longrightarrow A = D_n/C_n \simeq \mathbf{Z}/2\mathbf{Z}$ is the projection. (This provides the examples mentioned in Exercise 4.4.8.)

Exercise 4.6.6 (Two non-isomorphic groups with the same character table). Consider the dihedral group $G_1 = D_4$ of order 8, and the group G_2 defined as the subgroup of the multiplicative group of the Hamilton quaternions generated by i, j and k, or in other words (for readers unfamiliar with quaternions) the group generated by symbols i, j, k, subject to the relations

$$i^2 = j^2 = k^2 = ijk = -1.$$

(1) Show that G_2 is of order 8 (by enumerating its elements for instance), and that it contains a single element of order 2. Deduce that G_2 is not isomorphic to G_1.

(2) Compute the character table of G_2. Show that, up to possible reordering of the rows and columns, *it is identical with that of* G_1.

(3) Deduce from this a few things about a finite group that the character table *can not* determine (try to find as many things as possible that are different in G_1 and G_2; note that (1) already gives examples, and you should try to find others). For instance, can one determine all subgroups of a finite group from the character table, and not just the normal ones?

4.6.4. A linear group. The building blocks of all finite groups are, in some precise sense, the simple groups. We now consider the representations of the simplest type of group which is closely related to an infinite family of non-abelian simple groups: the linear groups $G_p = \mathrm{GL}_2(\mathbf{F}_p)$ for p prime (the simple groups in question are the quotients $\mathrm{PSL}_2(\mathbf{F}_q)$, for $q \notin \{2,3\}$). In contrast to the previous case, some of the irreducible representations that arise can not easily be described at the level of actual actions of G_p on specific vector spaces: they are identified as *characters*.

The whole computation is quite a bit more involved than in the previous cases, as can be expected, and the reader should be active in checking the details. For other fairly detailed accounts along the same lines, see [**13**, §5.2] or [**11**,], and for a treatment from a slightly different perspective, see [**7**, §5.1].

Before coming to this, we proceed in the usual way by finding the size of G_p and its conjugacy classes. For the first, the number of elements of G_p is the same as the number of bases of the plane $\mathbf{F}_p^2$, i.e.

$$|G_p| = (p^2-1)(p^2-p) = p(p-1)^2(p+1)$$

(there are p^2-1 choices for the first basis vector, and then the second may be any vector except the p which are linearly dependent on the first one).

We will assume that $p \geqslant 3$ (and briefly mention the simpler case of $\mathrm{GL}_2(\mathbf{F}_2)$ in a final remark). Determining the conjugacy classes is a question of linear algebra over $\mathbf{F}_p$, and we can argue from the characteristic polynomial of a given element $g \in G_p$:

- If g has a multiple eigenvalue in $\mathbf{F}_p$, but is diagonalizable, this means g is a scalar matrix. There are $p-1$ such matrices, and each is a conjugacy class of size 1, which together form the center Z of G_p;
- If g has a multiple eigenvalue but is *not* diagonalizable, we can find a basis of $\mathbf{F}_p^2$ in which g has the form

  $$g = \begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$$

  (first we can conjugate g to triangular form, but then the argument of the previous section gives a conjugate as above.) There are $p-1$ such conjugacy classes, and to compute their size we leave it to the reader to check that the centralizer of such a matrix g is the subgroup

  $$K = \left\{ \begin{pmatrix} x & t \\ 0 & x \end{pmatrix} \mid x \in \mathbf{F}_p^\times, t \in \mathbf{F}_p \right\}$$

  so that the size of the conjugacy class of g is $|B_p/K| = p^2-1$;
- If g has two distinct eigenvalues *in* $\mathbf{F}_p$, it is diagonalizable over $\mathbf{F}_p$, i.e., it is conjugate to a matrix

  $$g = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$$

with $a \neq b$. However, there are only $\frac{1}{2}(p-1)(p-2)$ such classes because one can permute a and b by conjugating with

$$w = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

and each of these classes has size $|G_p|/(p-1)^2 = p(p+1)$ because the centralizer of g as above is easily checked to be the group

$$T_1 = \Big\{ \begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix} \mid x,\ y \neq 0 \Big\} \tag{4.44}$$

of diagonal (not necessarily scalar) matrices, which is isomorphic obviously to $\mathbf{F}_p^\times \times \mathbf{F}_p^\times$, and is of order $(p-1)^2$;

- Finally, if g has two distinct eigenvalues, but they do not belong to the base field $\mathbf{F}_p$, the matrix can be diagonalized, but only over the extension field $k/\mathbf{F}_p$ generated by the eigenvalues. This is necessarily the unique (up to isomorphism) extension field of degree 2, and this is generated by some element α such that $\varepsilon = \alpha^2$ is a fixed *non-square* in $\mathbf{F}_p^\times$ (the existence of ε uses the assumption $p \geqslant 3$; for $p \equiv 3 \pmod 4$, for instance, one has to take $\varepsilon = -1$, though it is not in general possible to write down an exact formula for such an element). Once ε is fixed, we can see that g is conjugate to a matrix

$$g = \begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix} \in G_p$$

for some $a \in \mathbf{F}_p$ and $b \in \mathbf{F}_p^\times$. However, as before, the number of classes of this type is only $\frac{1}{2}p(p-1)$ (changing b into $-b$ does not change the conjugacy class.)

The centralizer of an element g of this type is seen to be equal to the subgroup

$$T_2 = \Big\{ \begin{pmatrix} x & y \\ \varepsilon y & x \end{pmatrix} \mid (x, y) \neq (0, 0) \Big\} \tag{4.45}$$

of order $p^2 - 1$, so that the conjugacy classes are of size $p(p-1)$. We observe that T_2 is in fact a rather simple group: it is abelian, and in fact the map

$$\begin{cases} T_2 & \longrightarrow & \mathbf{F}_p(\sqrt{\varepsilon}) \\ \begin{pmatrix} x & y \\ \varepsilon y & x \end{pmatrix} & \mapsto & x + y\sqrt{\varepsilon} \end{cases} \tag{4.46}$$

is an isomorphism. The determinant on T_2 corresponds, under this isomorphism to the map sending $x + y\sqrt{\varepsilon}$ to $x^2 - \varepsilon y^2$, which is the norm homomorphism $\mathbf{F}(\sqrt{\varepsilon})^\times \to \mathbf{F}_p^\times$ (this is because $\alpha^p = -\alpha$, since $\alpha = \sqrt{\varepsilon}$ generates the extension of $\mathbf{F}_p$ of degree 2).

Remark 4.6.7. Note the close formal similarity between the group T_2 and the group

$$\Big\{ \begin{pmatrix} x & y \\ -y & x \end{pmatrix} \mid x,\ y \in \mathbf{R},\ (x, y) \neq (0, 0) \Big\} \subset \mathrm{GL}_2(\mathbf{R})$$

which is isomorphic to $\mathbf{C}^\times$ by mapping an element as above to $x + iy$. Here, of course, the real number -1, which is not a square in $\mathbf{R}$, plays the role of ε. If $y \neq 0$, the corresponding matrix in $\mathrm{GL}_2(\mathbf{R})$ is not diagonalizable over $\mathbf{R}$, but it is over $\mathbf{C}$, with conjugate eigenvalues $x \pm iy$.

Tallying all this, we see that the number of conjugacy classes is

$$|G_p^\sharp| = 2(p-1) + \tfrac{1}{2}(p-1)(p-2) + \tfrac{1}{2}p(p-1) = p^2 - 1.$$

Remark 4.6.8. The following terminology is used for these four types of conjugacy classes: they are (1) scalar classes; (2) non-semisimple; (3) split semisimple; (4) non-split semisimple, respectively. The fourth type did not appear in the previous section, whereas the first three do intersect the upper-triangular subgroup B_p.

It will not be very difficult to find three "families" of irreducible representations. Once this is done, we will see what is missing.

First, it is well-known that the commutator group of G_p is $\mathrm{SL}_2(\mathbf{F}_p)$, and thus the determinant gives an isomorphism

$$\det\,:\,G_p/[G_p,G_p]\longrightarrow\mathbf{F}_p^{\times},$$

so that we have $p-1$ characters of dimension 1 given by

$$\chi(g)=\chi_1(\det(g))$$

for some character of $\mathbf{F}_p^{\times}$.

The next construction is based on induction: we use the subgroup B_p to induce its one-dimensional characters

$$\varrho(\chi_1,\chi_2)\,:\,\begin{pmatrix}a&t\\0&b\end{pmatrix}\mapsto\chi_1(a)\chi_2(b)$$

where χ_1 and χ_2 are again characters of $\mathbf{F}_p^{\times}$, and we denote

$$\pi(\chi_1,\chi_2)=\mathrm{Ind}_{B_p}^{G_p}(\varrho(\chi_1,\chi_2)),$$

which has dimension $p+1$. To compute its character, we use the set of representatives

$$R=\Bigl\{\begin{pmatrix}1&0\\t&1\end{pmatrix}\,,\;t\in\mathbf{F}_p,\begin{pmatrix}0&1\\1&0\end{pmatrix}\Bigr\}$$

of the cosets $B_p\backslash H_p$. Thus, for $\pi=\pi(\chi_1,\chi_2)$ we have

$$\chi_\pi(g)=\sum_{\substack{r\in R\\ rgr^{-1}\in B_p}}\chi_1(a_r)\chi_2(b_r)\tag{4.47}$$

(where we write a_r and b_r for the diagonal coefficients of rgr^{-1}). Before considering the four conjugacy types in turn, we observe the following very useful fact: for any $x\in G_p$, we have

$$xB_px^{-1}\cap B_p=\begin{cases}B_p\text{ if }x\in B_p\\ \{g\mid g\text{ diagonal in the basis }(e_1,xe_1)\}\text{ if }x\notin B_p,\end{cases}\tag{4.48}$$

Indeed, by definition, an element of B_p has the first basis vector $e_1\in\mathbf{F}_p^2$ as eigenvector, and an element of xB_px^{-1} has xe_1 as eigenvector; if xe_1 is not proportional to e_1 – i.e., if $x\notin B_p$ – this means that $g\in B_p\cap xB_px^{-1}$ if (and only if) g is diagonalizable in the fixed basis (e_1,xe_1). (Note that in this case, the intersection is a specific conjugate of the group T_1 of diagonal matrices.)

Now we compute:

- If $g=\begin{pmatrix}a&0\\0&a\end{pmatrix}$ is scalar, we obtain $\chi_\pi(g)=(p+1)\chi_1(a)\chi_2(a)$;

- If $g=\begin{pmatrix}a&1\\0&a\end{pmatrix}$ is not semisimple, since it is in B_p, only $1\in R$ contributes to the sum (since for $r\neq 1$ to contribute, it would be necessary that $g\in r^{-1}B_pr\cap B_p$, which is not possible by (4.48) since g is not diagonalizable). Thus we get

$$\chi_\pi(g)=\chi_1(a)\chi_2(a)$$

in that case;

- If $g=\begin{pmatrix}a&0\\0&b\end{pmatrix}$ with $a\neq b$, besides $r=1$, the other contributions must come from $r\in R$ such that g is diagonal in the basis (e_1,re_1), which is only possible if $re_1=e_2$, i.e., the only other possibility is $r=w$; this gives

$$\chi_\pi(g)=\chi_1(a)\chi_2(b)+\chi_1(b)\chi_2(a)$$

for the split semisimple elements;

- If g is non-split semisimple, it has no conjugate at all in B_p (as this would mean that g has an eigenvalue in $\mathbf{F}_p$), and hence $\chi_\pi(g)=0$.

The character values are therefore quite simple:

	$\begin{pmatrix}a&0\\0&a\end{pmatrix}$	$\begin{pmatrix}a&1\\0&a\end{pmatrix}$	$\begin{pmatrix}a&0\\0&b\end{pmatrix}$, $b\neq a$	$\begin{pmatrix}a&b\\\varepsilon b&a\end{pmatrix}$, $b\neq 0$
$\pi(\chi_1,\chi_2)$	$(p+1)\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(b)+\chi_2(b)\chi_1(a)$	0

Table 4.4. Character of $\pi(\chi_1,\chi_2)$

The squared norm of the character of these induced representations is also quite straightforward to compute: we find

$$\langle\chi_\pi,\chi_\pi\rangle=\frac{1}{|G_p|}\Big\{(p-1)(p+1)^2+(p-1)(p^2-1)+A\Big\}$$

where A is the contribution of the split semisimple classes, namely

$$A=\tfrac{1}{2}p(p+1)\sum\sum_{\substack{a,b\in\mathbf{F}_p^\times\\a\neq b}}|\chi_1(a)\chi_2(b)+\chi_1(b)\chi_2(a)|^2.$$

To compute A, one can expand the modulus squared, obtaining

$$A=\frac{p(p+1)}{2}\Big(2(p-1)(p-2)+2\,\mathrm{Re}(B)\Big)$$

with

$$\begin{aligned}B&=\sum\sum_{\substack{a,b\in\mathbf{F}_p^\times\\a\neq b}}\chi_1(a)\overline{\chi_1(b)\chi_2(a)}\chi_2(b)\\&=\sum\sum_{a,b}\chi_1(a)\overline{\chi_1(b)\chi_2(a)}\chi_2(b)-(p-1)\\&=\Big(\sum_{x\in\mathbf{F}_p^\times}\chi_1(x)\overline{\chi_2(x)}\Big)^2-(p-1).\end{aligned}$$

Thus there are two cases: if $\chi_1 = \chi_2$, we have $B = (p-1)^2 - (p-1)$, whereas if $\chi_1 \neq \chi_2$, we get $B = -(p-1)$. This leads, if no mistake is made in gathering all the terms, to

$$\langle \pi(\chi_1, \chi_2), \pi(\chi_1, \chi_2) \rangle = \begin{cases} 2 & \text{if } \chi_1 = \chi_2 \\ 1 & \text{if } \chi_1 \neq \chi_2. \end{cases} \tag{4.49}$$

Thus $\pi(\chi_1, \chi_2)$ is irreducible if and only $\chi_1 \neq \chi_2$ (see Exercise 4.8.3 for another argument towards this result, which is less computational). This means that we have found many irreducible representations of dimension $p+1$. the precise number is $\frac{1}{2}(p-1)(p-2)$, because in addition to requiring $\chi_1 \neq \chi_2$, we must remove the possible isomorphisms between those representations, and the character values show that if $\chi_1 \neq \chi_2$, we have

$$\pi(\chi_1, \chi_2) \simeq \pi(\chi_1', \chi_2') \text{ if and only if } (\chi_1', \chi_2') = (\chi_1, \chi_2) \text{ or } (\chi_1', \chi_2') = (\chi_2, \chi_1).$$

These $\frac{1}{2}(p-1)(p-2)$ representations are called the *principal series* representations for $\mathrm{GL}_2(\mathbf{F}_p)$.

Remark 4.6.9. The existence of the isomorphism $\pi(\chi_1, \chi_2) \simeq \pi(\chi_2, \chi_1)$ is guaranteed by the equality of characters. It is not immediate to *write down* an explicit isomorphism. (Note that, by Schur's Lemma, we have $\dim \mathrm{Hom}_{G_p}(\pi(\chi_1, \chi_2), \pi(\chi_2, \chi_1)) = 1$, so at least the isomorphism is unique, up to scalar; for an actual description, see, e.g., [**7**, p.404].)

Even when $\chi_1 = \chi_2$ we are not far from having an irreducible: since

$$\langle \pi(\chi_1, \chi_1), \pi(\chi_1, \chi_1) \rangle = 2,$$

the induced representation has two irreducible components. Could it be that one of them is one-dimensional? Using Frobenius reciprocity we see that for a one-dimensional character of the type $\chi \circ \det$, we have

$$\langle \pi(\chi_1, \chi_1), \chi \circ \det \rangle_{G_p} = \langle \varrho(\chi_1, \chi_1), \varrho(\chi, \chi) \rangle_{B_p} = \begin{cases} 0 & \text{if } \chi_1 \neq \chi \\ 1 & \text{if } \chi_1 = \chi, \end{cases}$$

since the restriction of $\chi \circ \det$ to B_p is

$$\begin{pmatrix} a & t \\ 0 & b \end{pmatrix} \mapsto \chi(a)\chi(b).$$

Switching notation, we see that $\pi(\chi, \chi)$ contains a unique 1-dimensional representation, which is $\chi \circ \det$. Its other component, denote $\mathrm{St}(\chi)$, is irreducible of dimension p, with character values given by

$$\chi_{\mathrm{St}(\chi)} = \chi_{\pi(\chi,\chi)} - \chi \circ \det,$$

namely

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, b \neq a$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}, b \neq 0$
$\mathrm{St}(\chi)$	$p\chi(a)^2$	0	$\chi(a)\chi(b)$	$-\chi(a^2 - \varepsilon b^2)$

From this, we see also that these representations are pairwise non-isomorphic (use the values for split semisimple elements).

Another description of these representations, which are called the *Steinberg* representations, is the following: first G_p acts on the set X_p of lines in $\mathbf{F}_p^2$, as did B_p in the previous section; by linear algebra, this action is doubly transitive (choosing two vectors on a pair of distinct lines gives a basis of $\mathbf{F}_p^2$, and any two bases can be mapped to one another using G_p), and therefore by Proposition 4.3.17, the permutation representation associated to X_p splits as

$$\mathbf{1} \oplus \mathrm{St}$$

for some irreducible representation St of dimension p. Then we get

$$\mathrm{St}(\chi) \simeq \mathrm{St} \otimes \chi(\det)$$

(e.g., because the permutation representation on X_p is isomorphic to $\mathrm{Ind}_{B_p}^{G_p}(\mathbf{1}) = \pi(\mathbf{1}, \mathbf{1})$, as in Example 2.6.4, (2), so that St is the same as $\mathrm{St}(\mathbf{1})$, and then one can use character values to check the effect of multiplying with $\chi \circ \det$.) The character of "the" Steinberg representation is particularly nice-looking:

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, b \neq a$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}, b \neq 0$
St	p	0	1	-1

To summarize our situation: we have found

$$p-1, \qquad \tfrac{1}{2}(p-1)(p-2), \qquad p-1$$

irreducible representations of dimension

$$1, \qquad p+1, \qquad p$$

respectively. There remains to find $\frac{1}{2}p(p-1)$ representations, with sum of squares of dimensions equal to

$$|G_p| - (p-1) - (p-1)p^2 - \tfrac{1}{2}(p-1)(p-2)(p+1)^2 = \tfrac{1}{2}p(p-1)(p-1)^2.$$

It seems therefore to be an excellent guess that the representations in question should be of dimension $p-1$... They will be the *discrete series* or *cuspidal* representations of G_p.

Note already the striking parallel between the (known) rows and columns of the evolving character table: for the first three families of conjugacy classes, we have found families of the same number of irreducible representations, all with a common dimension. We can therefore indeed expect to find a last family, which should correspond somehow to the non-split semisimple conjugacy classes of G_p. Another clear reason for the existence of a link with these conjugacy classes is that, for the moment, any combination of "known" irreducible characters is a function of $\det(g) = a^2 - \varepsilon b^2$ only when restricted to a non-split conjugacy class.

As it turns out, just as the induced representations $\pi(\chi_1, \chi_2)$ are parametrized by the pair (χ_1, χ_2), which can be interpreted as a character of the centralizer T_1 of a split semisimple conjugacy class (see (4.44)), the cuspidal representations a parametrized by certain characters ϕ of the common (abelian) centralizer T_2 of the representatives we use for the non-split conjugacy class, defined in (4.45). We pull the formula out of a hat, as a class function: we identify

$$\phi : T_2 \longrightarrow \mathbf{C}^\times$$

with a character $\mathbf{F}_p(\sqrt{\varepsilon})^\times \longrightarrow \mathbf{C}^\times$ using the isomorphism (4.46), and define a function $R(\phi)$ by

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$, $b \neq a$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}$, $b \neq 0$
$R(\phi)$	$(p-1)\phi(a)$	$-\phi(a)$	0	$-(\phi(a+b\sqrt{\varepsilon})+\phi(a-b\sqrt{\varepsilon}))$

We claim these give us the missing characters, for suitable ϕ.

Proposition 4.6.10. *Let ϕ be a character of T_2, or equivalently of $\mathbf{F}_p(\sqrt{\varepsilon})^\times$, such that $\phi \neq \phi'$, where the character ϕ' is defined by*[11]

$$\phi'(x+y\sqrt{\varepsilon}) = \phi(x-y\sqrt{\varepsilon}). \tag{4.50}$$

Then $R(\phi)$ is an irreducible character of G_p. Moreover, we have

$$R(\phi_1) = R(\phi_2)$$

if and only if either $\phi_1 = \phi_2$ or $\phi_1 = \phi_2'$.

Once this is known, we have all the characters we need. Indeed, these characters are of dimension $p-1$. To count them, we note that the condition (4.50) is equivalent with

$$\ker(\phi) \supset \{w \in \mathbf{F}_p(\sqrt{\varepsilon})^\times \mid w = \frac{x+y\sqrt{\varepsilon}}{x-y\sqrt{\varepsilon}}\}$$

(when seeing ϕ as a character of $\mathbf{F}_p(\sqrt{\varepsilon})^\times$) and the right-hand side is the same as the kernel of the norm map

$$\mathbf{F}_p(\sqrt{\varepsilon})^\times \longrightarrow \mathbf{F}_p^\times.$$

Thus those ϕ which *do* satisfy the condition are in bijection with the characters of the image of the norm map, which is $\mathbf{F}_p^\times$ since the norm is surjective. There are therefore $p-1$ (the characters of the form

$$x+y\sqrt{\varepsilon} \mapsto \chi(x^2-\varepsilon y^2)$$

where χ is a character of $\mathbf{F}_p^\times$) to be excluded from the p^2-1 characters of T_2. Finally, the identities $R(\phi) = R(\phi')$ show that the total number of irreducible characters given by the proposition is, as expected, $\frac{1}{2}p(p-1)$.

Proof of Proposition 4.6.10. We see first that the identity $R(\phi) = R(\phi')$ does hold, as equality of class functions. Similarly, the restriction $\phi \neq \phi'$ is a necessary condition for $R(\phi)$ to be an irreducible character, as we see by computing the square norm, which should be equal to 1: we have

$$\langle R(\phi), R(\phi)\rangle = \frac{1}{|G_p|}\Big\{(p-1)^3 + (p-1)(p^2-1) + \\ \frac{1}{2}p(p-1)\sum_{\substack{a,b\in\mathbf{F}_p \\ b\neq 0}} |\phi(a+b\sqrt{\varepsilon})+\phi(a-b\sqrt{\varepsilon})|^2\Big\}$$

[11] A better notation would be $\phi' = \phi^p$, since this is what the operation $a+b\sqrt{\varepsilon} \mapsto a-b\sqrt{\varepsilon}$ amounts to.

and the last sum (rather like the one for the induced representation $\pi(\chi_1, \chi_2)$) is equal to

$$\begin{aligned}\sum_{\substack{a,b\in\mathbf{F}_p\\ b\neq 0}} |\phi(a+b\sqrt{\varepsilon})+\phi(a-b\sqrt{\varepsilon})|^2 &= 2p(p-1)+2\,\mathrm{Re}\Bigl(\sum_{\substack{a,b\in\mathbf{F}_p\\ b\neq 0}} \phi(a+b\sqrt{\varepsilon})\overline{\phi(a-b\sqrt{\varepsilon})}\Bigr)\\ &= 2p(p-1)+2\,\mathrm{Re}\Bigl(\sum_{x\in\mathbf{F}_p(\sqrt{\varepsilon})^\times} \phi(x)\overline{\phi'(x)}-\sum_{a\in\mathbf{F}_p^\times} 1\Bigr)\\ &= 2(p-1)^2+(p^2-1)\langle\phi,\phi'\rangle\end{aligned}$$

where the last inner product can be seen on the group T_2. Thus we carefully find

$$\langle R(\phi), R(\phi)\rangle = 1+\langle\phi,\phi'\rangle,$$

which is 1 if and only if $\phi\neq\phi'$.

This result, and similar checks (one may verify in similar manner that $R(\phi)$, as a class function, is orthogonal to all the irreducible characters previously known), show that $R(\phi)$ behaves like the character of an irreducible representation. But this strong evidence is not, by itself, conclusive: although it shows that $R(\phi)$, when expanded into a combination of characters, must only involve the missing ones, this does not by itself guarantee that it is one itself.

This is something we noticed already in Remark 4.3.18; and as in that remark, we see at least (since $R(\phi)$ has norm 1 and $R(\phi)$ takes positive value at 1) that in order to conclude, it is enough to exhibit a linear combination of characters with *integral* coefficients which is equal to $R(\phi)$. This we do as in [**13**], though with even less motivation: we claim that

$$R(\phi) = \chi_1-\chi_2-\chi_3,$$

where χ_i is the character of the representation ϱ_i given by

$$\begin{aligned}\varrho_1 &= \pi(\phi,\mathbf{1})\otimes\mathrm{St}(\phi), \quad \text{(where } \phi \text{ is restricted to } \mathbf{F}_p^\times)\\ \varrho_2 &= \pi(\phi,\mathbf{1}),\\ \varrho_3 &= \mathrm{Ind}_{T_2}^{G_p}(\phi).\end{aligned}$$

Checking this is a matter of computation; note at least that the dimension

$$p(p+1)-(p+1)-[G_p:T_2] = p^2-1-p(p-1) = p-1$$

is correct; the reader should of course make sure of the other values; we only give the character of the induced representation $\mathrm{Ind}_{T_2}^{G_p}(\phi)$ to facilitate the check if needed:

	$\begin{pmatrix} a & 0\\ 0 & a\end{pmatrix}$	$\begin{pmatrix} a & 1\\ 0 & a\end{pmatrix}$	$\begin{pmatrix} a & 0\\ 0 & b\end{pmatrix}$, $b\neq a$	$\begin{pmatrix} a & b\\ \varepsilon b & a\end{pmatrix}$, $b\neq 0$
$\mathrm{Ind}_{T_2}^{G_p}(\phi)$	$p(p-1)\phi(a)$	0	0	$\phi(a+b\sqrt{\varepsilon})+\phi(a-b\sqrt{\varepsilon})$

(this is especially easy to evaluate because T_2 only intersects conjugacy classes of central and non-split semisimple elements.) □

We are thus done computing this character table! To summarize, we present it in a single location:

	$\begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$	$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}, b \neq a$	$\begin{pmatrix} a & b \\ \varepsilon b & a \end{pmatrix}, b \neq 0$
$\chi \circ \det$	$\chi(a^2)$	$\chi(a^2)$	$\chi(ab)$	$\chi(a^2 - \varepsilon b^2)$
$\pi(\chi_1, \chi_2)$	$(p+1)\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(a)$	$\chi_1(a)\chi_2(b) + \chi_2(b)\chi_1(a)$	0
$\mathrm{St}(\chi)$	$p\chi(a^2)$	0	$\chi(ab)$	$-\chi(a^2 - \varepsilon b^2)$
$R(\phi)$	$(p-1)\phi(a)$	$-\phi(a)$	0	$-(\phi(a + b\sqrt{\varepsilon})+$ $\phi(a - b\sqrt{\varepsilon}))$

Table 4.5. Character table of $\mathrm{GL}_2(\mathbf{F}_p)$

Remark 4.6.11. (1) For $p = 2$, the only actual difference is that there are no split semisimple conjugacy classes, and correspondingly no principal series (induced) representations. Indeed, $\mathrm{GL}_2(\mathbf{F}_2)$ is isomorphic to $\mathfrak{S}_3$ (an isomorphism is obtained by looking at the permutations of the three lines in $\mathbf{F}_2^2$ induced by an element of $\mathrm{GL}_2(\mathbf{F}_2)$), and the character table of the latter in Example 4.6.1 corresponds to the one above when we remove the third line and column: the 2-dimensional representation of $\mathfrak{S}_3$ corresponds to the (unique) Steinberg representation and the signature corresponds to the (unique) cuspidal representation of $\mathrm{GL}_2(\mathbf{F}_2)$.

(2) The restriction to $\mathrm{GL}_2(k)$ where k is a field of prime order was merely for convenience; all the above, and in particular the full character table, are valid for an arbitrary finite field k, with characters of $k^\times$, and of the group of invertible elements in its quadratic extension, instead of those of $\mathbf{F}_p^\times$ and $\mathbf{F}_{p^2}^\times$.

This computation of the character table of $\mathrm{GL}_2(\mathbf{F}_q)$ was somewhat involved. The following series of exercises shows some of the things that can be done once it is known.

Exercise 4.6.12 (Characters of $\mathrm{SL}_2(\mathbf{F}_p)$). The group $\mathrm{SL}_2(\mathbf{F}_p)$ is quite closely related to $\mathrm{GL}_2(\mathbf{F}_p)$, and one can compute the character table of one from that of the other.

(1) For $p \geqslant 3$, show that $\mathrm{SL}_2(\mathbf{F}_p)$ has $p + 4$ conjugacy classes, and describe representatives of them.

(2) By decomposing the restriction to $\mathrm{SL}_2(\mathbf{F}_p)$ of the irreducible representations of $\mathrm{GL}_2(\mathbf{F}_p)$, describe the character table of $\mathrm{SL}_2(\mathbf{F}_p)$ for $p \geqslant 3$. [See, for instance, [**13**, §5.2] for the results; there are two irreducible representations of $\mathrm{GL}_2(\mathbf{F}_p)$ which decompose as a direct sum of two representations whose characters are quite tricky to compute, and you may try at first to just compute the dimensions of the irreducible components.]

(3) Show in particular that

$$\min_{\pi \neq \mathbf{1}} \dim \pi = \frac{p-1}{2} \tag{4.51}$$

where π runs over all non-trivial irreducible (complex) representations of $\mathrm{SL}_2(\mathbf{F}_p)$.

(In Section 4.7.1, we will see some striking applications of the fact that this dimension is large, in particular that it tends to infinity as p does, and we will prove (4.51) more directly, independently of the computation of the full character table.)

(4) For $G = \mathrm{GL}_2(\mathbf{F}_p)$, $H = \mathrm{SL}_2(\mathbf{F}_p)$, $A = G/H \simeq \mathbf{F}_p^\times$, compute the invariant $\kappa(\pi)$ defined in (4.34) for all representations π of G.

Exercise 4.6.13 (The Gelfand-Graev representation). Let

$$U = \left\{ \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \mid t \in \mathbf{F}_p \right\} \subset G_p,$$

This is a subgroup of G_p, isomorphic to $\mathbf{F}_p$. Let $\psi \neq 1$ be a non-trivial irreducible character of U. Compute the character of $\varrho = \mathrm{Ind}_U^{G_p}(\psi)$ and show that if π is an irreducible representation of G_p, we have

$$\langle \varrho, \pi \rangle = \begin{cases} 1 & \text{if } \dim(\pi) \geqslant 2, \\ 0 & \text{otherwise.} \end{cases}$$

The representation ϱ is called the *Gelfand-Graev* representation of $\mathrm{GL}_2(\mathbf{F}_p)$. In concrete terms, the result means that for any irreducible representation

$$\pi : G_p \longrightarrow \mathrm{GL}(E),$$

which is not one-dimensional, there exists, up to scalar, a unique linear form (called a *Whittaker functional* for π)

$$\ell_\pi : E \longrightarrow \mathbf{C}$$

such that

$$\ell_\pi\Bigl(\pi\begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} v\Bigr) = \psi(x)\ell_\pi(v)$$

for $x \in \mathbf{F}_p$ and $v \in E$ (this is because such a linear form is exactly an element of $\mathrm{Hom}_U(\pi, \psi)$, which is isomorphic to $\mathrm{Hom}_G(\pi, \varrho)$ by Frobenius reciprocity).

Using the specific isomorphism that implements Frobenius reciprocity, we find that given such a linear form $\ell_\pi \neq 0$, the homomorphism

$$\pi \longrightarrow \mathrm{Ind}_U^G(\psi)$$

is given by mapping a vector v to the function

$$W_v(g) = \ell_\pi(\varrho(g)v).$$

Exercise 4.6.14 (Distinct characters that coincide on a generating set). Show that the following can happen for some finite groups: there may exist a group G, a *generating set* S, and two irreducible (even faithful) representations ϱ_1 and ϱ_2 which are non-isomorphic but satisfy

$$\chi_{\varrho_1}(s) = \chi_{\varrho_2}(s)$$

for all $s \in S$.

Exercise 4.6.15. Let ϱ be any irreducible complex representation of $\mathrm{GL}_2(\mathbf{F}_p)$. Let χ be the character of $\mathbf{F}_p^\times$ such that

$$\varrho\Bigl(\begin{pmatrix} x & 0 \\ 0 & x \end{pmatrix}\Bigr) = \chi(x)\mathrm{Id}$$

(the central character of ϱ). Show that the contragredient $\tilde{\varrho}$ of ϱ is isomorphic to the representation $\hat{\varrho}$ given by

$$\hat{\varrho}(g) = \chi(\det(g))^{-1}\varrho(g).$$

[Hint: This can be done without using the character table, by looking at what is the transpose of g^{-1}.]

Exercise 4.6.16 (Commutators in $\mathrm{GL}_2(\mathbf{F}_p)$). Using Proposition 4.4.3, show that the *set* of commutators (not the subgroup they generate!) in $\mathrm{GL}_2(\mathbf{F}_p)$ is equal to $\mathrm{SL}_2(\mathbf{F}_p)$ for $p \geqslant 3$.

Exercise 4.6.17. (1) For $p \geqslant 3$ and π an irreducible representation of $\mathrm{GL}_2(\mathbf{F}_p)$, show that there exists a constant $c_\pi \in \mathbf{C}$ and a character χ of $\mathbf{F}_p^\times$ such that

$$\chi_\pi\Big(\begin{pmatrix} x & 1 \\ & x \end{pmatrix}\Big) = c_\pi \chi(x)$$

for all $x \in \mathbf{F}_p^\times$ (this can be done without the character table).

(2) Let f denote the characteristic function of the set of all $g \in \mathrm{GL}_2(\mathbf{F}_p)$ which are diagonalizable over an algebraic closure of $\mathbf{F}_p$. Show that

$$\langle f, \chi_\pi \rangle = 0$$

for all except p irreducible representations of $\mathrm{GL}_2(\mathbf{F}_p)$. [Hint: Here you should probably use the character table.]

4.6.5. The symmetric groups. The irreducible representations of the symmetric groups $\mathfrak{S}_n$, $n \geqslant 1$, were already determined essentially by Frobenius. Since then, there have been many different interpretations or variants of the construction and the subject remains a very lively topic of current research, both for its own sake, or because of its many applications.

We will give a short description of the irreducible representations, in the language of "Specht modules", but we will not given full proofs – there are many detailed treatments in the literature, including those in [**13**, Ch. 4] or [**8**, §28] and the very concise version in [**9**, Ch. 7].

4.7. Applications

We present in this section some sample applications of the representation theory of finite groups, where the statements do not, by themselves, seem to depend on representations. The sections are independent of each other.

4.7.1. "Quasirandom" groups. Quite recently, Gowers [**16**] introduced a notion of "quasirandom" groups, motivated in part by similar ideas in the context of graph theory. Here is one of the simplest results that can be obtained in this area:

Theorem 4.7.1 (Product decompositions with small sets; Gowers, Nikolov-Pyber). *Let $G \neq 1$ be a non-trivial finite group and $k \geqslant 1$ the smallest dimension of a non-trivial irreducible complex representation of G. For any subsets A, B, C in G such that*

$$\frac{|A||B||C|}{|G|^3} > \frac{1}{k}, \tag{4.52}$$

we have $ABC = G$, or in other words, every *element $g \in G$ can be written as $g = abc$ with $a \in A$, $b \in B$ and $c \in C$.*

We use here the following product notation: for subsets A_1, ..., A_k of a group G (not necessarily distinct), the set $A_1A_2\cdots A_k$ is the set of all products

$$a = a_1a_2\cdots a_k$$

with $a_i \in A_i$ for all i; if some $A_i = \{a_i\}$ are singletons, we can just write the corresponding element a_i, e.g., in $a_1A_2a_3$. We also write A_1^{-1} for the set of all a^{-1} with $a \in A$.

It is also convenient to denote

$$\nu(A) = \frac{|A|}{|G|}$$

for $A \subset G$: this is the "density" of A in G, It can be interpreted intuitively as the probability that a "random element" in G belong to A, and the hypothesis (4.52) of the theorem can be phrased as

$$\nu(A)\nu(B)\nu(C) > \frac{1}{k}. \tag{4.53}$$

Proof. The first step is due to Gowers [**16**, Lemma 5.1]: under the stated condition (4.53), we will show that $AB \cap C$ is not empty, i.e., that some $c \in C$ is of the form ab with $a \in A$ and $b \in B$.

To proceed with better motivation, fix only the two sets B and C. We try to find an upper bound on the size of the set D of those elements $g \in G$ such that the intersection

$$C \cap gB$$

is empty; indeed, to say that a set A fails to satisfy $AB \cap C \neq \emptyset$ is to say that $A \subset D$, and if we know that D has a certain size, then it can not contain any set of larger size.

The idea to control $|D|$ (or the density $\nu(D)$) is to look at the function

$$\varphi_{B,C} \,:\, g \mapsto \frac{|C \cap gB|}{|G|} = \nu(C \cap gB)$$

defined on G, and to show that it is non-zero on a relatively large set by finding an upper bound for its "variance", i.e., the mean-square of $\varphi_{B,C}$ minus its average.

This average value is easy to determine: we have

$$\langle \varphi_{B,C}, 1 \rangle = \frac{1}{|G|^2} \sum_{g \in G} |C \cap gB| = \frac{1}{|G|^2} \sum_{c \in C} \sum_{\substack{g \in G \\ c \in gB}} 1 = \frac{|B||C|}{|G|^2} = \nu(B)\nu(C),$$

since $c \in gB$ is equivalent with $g \in cB^{-1}$, which has order $|B|$. Hence we wish to understand the quantity

$$\frac{1}{|G|} \sum_{a \in G} \Bigl(\varphi_{B,C}(a) - \nu(B)\nu(C) \Bigr)^2,$$

and if we know an upper-bound (say V) for it, we can argue by positivity[12] that the set X of those $g \in G$ with $\varphi_{B,C} = 0$ (i.e., $C \cap gB = \emptyset$) satisfies

$$\frac{|X|}{|G|} (\nu(B)\nu(C))^2 = \frac{1}{|G|} \sum_{g \in X} \Bigl(\varphi_{B,C}(g) - \nu(B)\nu(C) \Bigr)^2 \leqslant V,$$

and in particular, if $A \subset G$ satisfies

$$\nu(A) > \frac{V}{(\nu(B)\nu(C))^2},$$

it must be the case that $\varphi_{B,C}$ is not identically zero on A, i.e., that $AB \cap C$ is not empty.

Now, in order to analyze $\varphi_{B,C}$, we observe that for any $g \in G$, we have

$$\varphi_{B,C}(g) = \frac{1}{|G|} \sum_{x \in G} \delta_C(x) \delta_{gB}(x) = \frac{1}{|G|} \sum_{x \in G} \delta_C(x) \delta_B(g^{-1}x)$$

where, for any subset $D \subset G$, we define

$$\delta_D(x) = \begin{cases} 1 & \text{if } x \in D, \\ 0 & \text{otherwise.} \end{cases}$$

[12] This is the trick known as Chebychev's inequality in probability theory.

In other words, defining $\psi(g) = \varphi_{B,C}(g^{-1})$, we have

$$\psi = \mathrm{reg}(\Delta_C)\delta_B,$$

where

$$\Delta_C = \frac{1}{|G|}\sum_{g\in G}\delta_C(g)g \in \mathbf{C}(G).$$

We now normalize ψ by subtracting the average, defining

$$\psi_0 = \psi - \langle\psi, 1\rangle = \mathrm{reg}(\Delta_C)\mu_B, \qquad \mu_B = \delta_B - \nu(B).$$

Our goal is then to bound from above the quantity

$$\langle\psi_0, \psi_0\rangle = \frac{1}{|G|}\sum_{g\in G}\Bigl(\psi(g) - \nu(B)\nu(C)\Bigr)^2 = \frac{1}{|G|}\sum_{a\in G}\Bigl(\varphi_{B,C}(a) - \nu(B)\nu(C)\Bigr)^2.$$

We do this by observing that $\mathrm{reg}(\Delta_C)$ is a linear map acting on the subspace

$$C_0(G) = \{\varphi \in C(G) \mid \langle\varphi, 1\rangle = 0\} \subset C(G)$$

and hence, by elementary Hilbert space theory, we have

$$\langle\psi_0, \psi_0\rangle \leqslant \lambda^2\langle\mu_B, \mu_B\rangle \tag{4.54}$$

where $\lambda^2 \geqslant 0$ is the largest eigenvalue of the non-negative self-adjoint operator

$$\Delta_2 = \mathrm{reg}(\Delta_C)^*\,\mathrm{reg}(\Delta_C)$$

acting on $C_0(G)$, the adjoint $\mathrm{reg}(\Delta_C)^*$ being computed for the inner product on $C(G)$.

We have not yet really used much representation theory. But here is the crux: consider the λ-eigenspace of Δ_2, say $E \subset C_0(G)$. Then E is a subrepresentation of the "left" regular representation, i.e., it is stable under the action of G such that

$$\mathrm{reg}'(g)\varphi(x) = \varphi(x^{-1}g),$$

simply because the two actions of G on itself by right and left multiplication commute: since $\mathrm{reg}(\Delta_C)$ is defined using right-multiplication, the operators $\mathrm{reg}'(g)$ commute with $\mathrm{reg}(\Delta_C)$ and its adjoint, hence with Δ_2, and therefore stabilize its eigenspaces. Indeed, if $\Delta_2\varphi = \lambda\varphi$, we have

$$\Delta_2(\mathrm{reg}'(g)\varphi) = \mathrm{reg}'(g)\,\mathrm{reg}(\Delta_2)\varphi = \lambda\,\mathrm{reg}'(g)\varphi.$$

Now our assumption shows that $\dim(E) \geqslant k$, because under reg', the invariant subspace of $C(G)$ is the space of constant functions, which is orthogonal to $C_0(G)$, so that reg' can not act trivially on any subspace of E. Thus the eigenvalues of Δ_2 have "large" multiplicity (if k is large).

How can this knowledge of the dimension of the eigenspace help bounding the eigenvalue? The point is that we can achieve some control of *all* the eigenvalues of Δ_2 using its trace, and because all eigenvalues are non-negative, we have

$$k\lambda^2 \leqslant (\dim E)\lambda^2 \leqslant \mathrm{Tr}(\Delta_2),$$

which we compute separately, using the relation

$$\mathrm{reg}(g)^* = \mathrm{reg}(g^{-1})$$

coming from unitarity, to obtain

$$\mathrm{Tr}(\Delta_2) = \frac{1}{|G|^2}\sum_{x,y\in G}\delta_C(x)\delta_C(y^{-1})\,\mathrm{Tr}(\mathrm{reg}(y^{-1}x)),$$

so that, by the character formula for the regular representation, we obtain

$$\mathrm{Tr}(\Delta_2) = \frac{1}{|G|} \sum_{\substack{x,y\in C\\ x=y}} 1 = \frac{|C|}{|G|} = \nu(C).$$

Thus we find an upper bound for λ^2, namely

$$\lambda^2 \leqslant \frac{\nu(C)}{k},$$

and hence by (4.54) we get

$$\frac{1}{|G|}\sum_{g\in G}\Bigl(\varphi_{B,C}(g) - \nu(B)\nu(C)\Bigr)^2 \leqslant \frac{\nu(C)}{k}\langle \mu_B, \mu_B\rangle.$$

But the last term is also easy to compute: we have

$$\begin{aligned}\langle \mu_B, \mu_B\rangle &= \langle \delta_B, \delta_B\rangle - 2\nu(B)\langle \delta_B, 1\rangle + \nu(B)^2\\ &= \nu(B)(1-\nu(B)) \leqslant \nu(B),\end{aligned}$$

and therefore[13] the conclusion is

$$\frac{1}{|G|}\sum_{g\in G}\Bigl(\varphi_{B,C}(g) - \nu(B)\nu(C)\Bigr)^2 \leqslant \frac{\nu(B)\nu(C)}{k}.$$

Now the positivity argument shows that the number, say N, of those $g \in G$ with $C \cap gB = \emptyset$ satisfies

$$\nu(N) \leqslant \frac{1}{k\nu(B)\nu(C)}.$$

This gives the intermediate statement proved by Gowers: if A, B, C satisfy

$$\frac{|A||B||C|}{|G|^3} = \nu(A)\nu(B)\nu(C) > \frac{1}{k},$$

the intersection $C \cap AB$ is not empty. Now we bootstrap this by an amazingly clever trick of Nikolov and Pyber [**30**, Prop. 1]: consider again A, B, C as in the proposition, and *redefine*

$$C_1 = G - AB = \{g \in G \mid g \text{ is not of the form } ab \text{ with } a \in A,\ b \in B\}.$$

Then *by definition*, we have $C_1 \cap AB = \emptyset$. By contraposition, using the result of Gowers, this means that we must have

$$\frac{|A||B||C_1|}{|G|^3} \leqslant \frac{1}{k}$$

and the assumption (4.53) now leads to

$$|C_1| < |C|.$$

This means that $|AB| + |C| > |G|$. Now for any $g \in G$, this means also that $|AB| + |gC^{-1}| = |AB| + |C| > |G|$. Therefore the sets AB and gC^{-1} must intersect; this precisely means that $g \in ABC$, and we are done. □

The next corollary does not mention representations at all:

[13] We could have kept the term $-\nu(B)^2$ to very slightly improve this estimate, but it does not seem to matter in any application.

Corollary 4.7.2 ($\mathrm{SL}_2(\mathbf{F}_p)$ is quasirandom). *If $p \geqslant 3$ is a prime number and $A \subset \mathrm{SL}_2(\mathbf{F}_p)$ is a subset such that*

$$|A| > 2^{1/3} p(p+1)(p-1)^{2/3}$$

or equivalently

$$\frac{|A|}{|\mathrm{SL}_2(\mathbf{F}_p)|} > \Bigl(\frac{2}{p-1}\Bigr)^{1/3},$$

then for any $g \in \mathrm{SL}_2(\mathbf{F}_p)$, there exist a_1, a_2, $a_3 \in A$ with $g = a_1 a_2 a_3$.

Proof. This follows from the proposition for $G = \mathrm{SL}_2(\mathbf{F}_p)$, where $|G| = p(p^2 + p)$, and $B = C = A$, using (4.51), which shows that $k = \frac{1}{2}(p-1)$. □

In particular, such sets are generators, but in a very strong sense. Results like this can be used to help with certain proofs of Theorem 1.2.5 in Section 1.2: to show that a very small subset like $S = \{s_1, s_2\}$ (as in (1.1)) generates $\mathrm{SL}_2(\mathbf{F}_p)$ in at worse $C \log p$ steps, it suffices to find some C' such that the number of elements in $\mathrm{SL}_2(\mathbf{F}_p)$ obtained using products of $\leqslant C' \log p$ elements from S is at least $2(p+1)^{8/9}$, for instance. Indeed, if that is the case, the set A formed by these products satisfies the assumption of the corollary, and every element of G is the product of at most $3C' \log p$ elements from S.

Exercise 4.7.3 (Minimal dimension of non-trivial representations of $\mathrm{SL}_2(\mathbf{F}_p)$). We indicate here how to prove (4.51) without invoking the full computation of the character table of $\mathrm{SL}_2(\mathbf{F}_p)$. Thus let $\varrho \neq \mathbf{1}$ be an irreducible unitary representation of $\mathrm{SL}_2(\mathbf{F}_p)$, with $p \geqslant 3$.

(1) Show that

$$A = \varrho\Bigl(\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}\Bigr)$$

is not the identity. [Hint: Use the fact that $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ generate $\mathrm{SL}_2(\mathbf{F}_p)$.]

(2) Let $\xi \in \mathbf{C}^\times$ be an eigenvalue of A with $\xi \neq 1$. Show that, for all a coprime to p, ξ^{a^2} is also an eigenvalue of A. [Hint: Use a suitable conjugate of A.]

(3) Deduce that $\dim(\varrho) \geqslant (p-1)/2$.

(Thanks to O. Dinai for pointing out this proof; note that it is only by constructing the "cuspidal" representations that it is possible to show that this bound is sharp, and also that if $\mathbf{F}_p$ is replaced with another finite field with q elements, of characteristic p, this argument does not give the correct lower bound $\frac{1}{2}(q-1)$.)

The terminology "quasirandom" may seem mysterious at first, but it is well explained by the mechanism of the proof: the average of the function $\varphi_{B,C}$ corresponds precisely to the intuitive "probability" that an element $x \in G$ belongs to two subsets of density $|B|/|G|$ and $|C|/|G|$ if these are genuinely random and independent. Hence, the fact that $\varphi_{B,C}$ is quite closely concentrated around its average value, when k is large, may be interpreted as saying that its elements and subsets behave as if they were random (in certain circumstances).

To put the result in context, note that if $k = 1$ (for instance if G is abelian, or if $G = \mathrm{GL}_2(\mathbf{F}_p)$, which has many one-dimensional irreducible representations) the condition (4.53) can not be satisfied unless $A = B = C = G$. And indeed a statement like Corollary 4.7.2 is completely false if, say, $G = \mathbf{Z}/p\mathbf{Z}$ with p large: for instance, if $A = B = C$ is the image modulo p of the set of integers $1 \leqslant n \leqslant \lfloor \frac{p}{3} \rfloor - 1$, we see that $A + B + C$ is not all of G, although the density of A is about $1/3$, for all p.

4.7.2. Burnside's "two primes" theorem. We prove here the theorem of Burnside mentioned in Chapter 1 (Theorem 1.2.7): a finite group with order divisible by at most two distinct primes is necessarily solvable. The proof is remarkable, in that it does not depend on being able to write the character table of the group being investigated, but on subtler features about a finite group that may be found by looking at its irreducible characters. These are related to *integrality properties*, which have many other important applications.

The basic idea is to prove the following, weaker-looking statement:

Proposition 4.7.4 (Existence of normal subgroup). *Let G be a finite group of order p^aq^b for some primes p and q and integers a, $b \geqslant 0$. If G is not abelian, it contains a normal subgroup $H \triangleleft G$ with $H \neq 1$ and $H \neq G$.*

To see that this implies Burnside's Theorem, that groups of order p^aq^b are solvable, one argues by induction on the order of a group G of this type. By the proposition, either G is abelian, and therefore solvable, or there exists $H \triangleleft G$ such that $H \neq 1$ and $G/H \neq 1$; in that case we have an exact sequence

$$1 \longrightarrow H \longrightarrow G \longrightarrow G/H \longrightarrow 1,$$

and both H and G/H have orders strictly smaller than $|G|$, and divisible only (at most) by the primes p and q. By induction, they are therefore solvable, and this is well-known to imply that G itself is solvable.

So we are reduced to a question of finding a non-trivial normal subgroup in a group G, one way or another, and Burnside's idea is to find it as the kernel of some suitable non-trivial irreducible representation

$$\varrho : G \longrightarrow \mathrm{GL}(E),$$

or of an associated homomorphism

$$\bar{\varrho} : G \xrightarrow{\varrho} \mathrm{GL}(E) \longrightarrow \mathrm{PGL}(E) ;$$

indeed, it is a bit easier to ensure that $\ker \bar{\varrho}$ is non-trivial (the kernel is enlarged modulo the scalars), and the possibility that $\ker \bar{\varrho} = G$ is so special that its analysis is even simpler.

We will find the desired representation by means of the following result, which is itself of great interest:

Theorem 4.7.5 (Burnside). *Let G be a finite group,*[14] *and*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

an irreducible complex representation of G. If $g \in G$ is such that its conjugacy class $g^\sharp \subset G$ has order coprime with $\dim \varrho$*, then either $\chi_\varrho(g) = 0$, or $g \in \ker \bar{\varrho}$, where $\bar{\varrho}$ is the composite homomorphism*

$$G \xrightarrow{\varrho} \mathrm{GL}(E) \longrightarrow \mathrm{PGL}(E).$$

This may not be a result that is easy to guess or motivate, except that the statement may well come to mind after looking at many examples of characters tables – for instance, in the case of the solvable groups B_p of order $p(p-1)^2$ (see Table 4.3 in Section 4.6.3), the characters of the irreducible representations of dimension $p-1$ vanish at all conjugacy classes of size $p-1$, and their values at conjugacy classes of size 1 are scalar matrices (hence in the kernel of $\bar{\varrho}$). Similarly, the character of the Steinberg representation of

[14] Of any order.

dimension p of $\mathrm{GL}_2(\mathbf{F}_p)$ is zero at all conjugacy classes of size p^2-1. (Note that if the reader did look, she will certainly have also remarked a striking fact: for any irreducible representation ϱ of dimension $\dim\varrho > 1$, there exists (or so it seems) some conjugacy class c with $\chi_\varrho(c) = 0$; this is indeed true, as we will explain in Remark 4.7.10 below...)

We can also check immediately that the statement of the theorem is true for conjugacy classes of size 1: this corresponds to elements of the center of G, for which $\varrho(g)$ is always a homothety for any irreducible representation ϱ (the central character, as in Corollary 2.7.15.)

Proof of Proposition 4.7.4 from Theorem 4.7.5. Note first that we can certainly assume that $a \geqslant 1$ and $b \geqslant 1$, since a group of order a power of a single prime has a non-trivial center (see, e.g., [**33**, Th. 4.4] for this basic feature of finite groups.)

We attempt to find an element $g \in G$ and an irreducible representation ϱ so that Theorem 4.7.5 applies, while ensuring that the character value $\chi_\varrho(g)$ is non-zero. The difficulty is to ensure the coprimality condition of $\dim(\varrho)$ with $|g^\sharp|$, and indeed this is where the assumption that $|G|$ is only divisible by two primes is important.

In fact, the following fact remains valid for arbitrary finite groups, and can be interpreted as one more attempt of conjugacy classes and irreducible representations to behave "dually" (see Remark 4.2.6):

Fact. Let $G \neq 1$ be a finite group, and let p, q be prime numbers. There exists a pair (g, ϱ), where $g \neq 1$ is an element of G, and $\varrho \neq \mathbf{1}$ is an irreducible complex representation of G, such that $\chi_\varrho(g) \neq 0$ and

$$p \nmid |g^\sharp|, \qquad q \nmid \dim(\varrho). \tag{4.55}$$

We first conclude the proof using this fact: the point is that if $|G| = p^a q^b$ with a, $b \geqslant 1$, then with g and ϱ as so conveniently given, the conditions (4.55) mean that $|g^\sharp|$ and $\dim(\varrho)$ must be coprime (one does not need to know the – true – fact that $\dim(\varrho) \mid |G|$: the order of $g^\sharp$ does divide $|G|$, and hence must be a power of q, but q is coprime with $\dim(\varrho)$).[15] Hence we can apply Theorem 4.7.5 and conclude that $g \in \ker\bar{\varrho}$, so that the latter is *non-trivial* normal subgroup. The endgame is now straightforward: if $\ker\bar{\varrho} \neq G$, this kernel is the required proper, non-trivial, normal subgroup, while otherwise the composition $\bar{\varrho}$ is trivial, and then ϱ takes scalar values, and must therefore be one-dimensional by irreducibility. We then get an isomorphism

$$G/\ker(\varrho) \simeq \mathrm{Im}(\varrho) \subset \mathbf{C}^\times,$$

which shows in turn that $\ker\varrho$ is a proper, non-trivial, normal subgroup... unless $G \simeq \mathrm{Im}(\varrho)$ is in fact abelian!

Now we prove the existence of the required pair (g, ϱ). Given $g \neq 1$, the basic relation between the values of the irreducible characters at g and their dimensions is the orthogonality relation (4.28), which gives

$$\sum_{\varrho\in\hat{G}} \chi_\varrho(g)\overline{\chi_\varrho(1)} = \sum_{\varrho\in\hat{G}} (\dim\varrho)\chi_\varrho(g) = 0.$$

[15] Of course, (4.55) does not exclude possibilities like

$$|G| = pqr, \quad g^\sharp = qr, \quad \dim(\varrho) = pr,$$

where p, q, r are distinct primes; see Remark 4.7.6 for the case of the alternating group A_5.

If we isolate, as usual, the contribution of the trivial representation, we find

$$\sum_{\varrho\neq\mathbf{1}}(\dim\varrho)\chi_\varrho(g)=-1, \tag{4.56}$$

which certainly tells us that there is some irreducible representation $\varrho\neq\mathbf{1}$ such that $\chi_\varrho(g)\neq 0$.

But even better, if we consider this modulo the prime number q, it implies that there is some irreducible representation $\varrho\neq\mathbf{1}$ with $\chi_\varrho(g)\neq 0$ and $q\nmid\dim\varrho$. This relies on the fact that the values $\chi_\varrho(g)$ of irreducible characters, which are sums of roots of unity (the eigenvalues of $\varrho(g)$) are *algebraic integers*: modulo ℓ, the right-hand side of (4.56) is non-zero, and some term in the sum is therefore not divisible by q.

Precisely, if it were the case that $q\mid\dim\varrho$ for all ϱ such that $\chi_\varrho(g)\neq 0$, we would get

$$-\frac{1}{q}=\sum_{\substack{\varrho\neq\mathbf{1}\\ q\mid\dim(\varrho)}}\Bigl(\frac{\dim(\varrho)}{q}\Bigr)\chi_\varrho(g), \tag{4.57}$$

where the right-hand side is an algebraic integer, and this is impossible since $1/q$ is not. In Section A.1 in the Appendix, we present a short discussion of the properties of algebraic integers that we use (here, Proposition A.1.1), and readers for whom this is not familiar may either read this now, or continue while assuming that the character values involved are all actual integers in $\mathbf{Z}$, since in that case (4.57) is patently absurd.

So, given $g\neq 1$ and the prime q, we can always find $\varrho\neq\mathbf{1}$ such that $q\nmid\dim\varrho$ and $\chi_\varrho(g)\neq 0$. It is therefore sufficient, to prove the claim, to show that $g\neq 1$ with $p\nmid|g^\sharp|$ also exists. Of course if $p\nmid|G|$, this is always true, and we can assume that p is a divisor of $|G|$. Then we use another "averaging" trick: by partitioning G into conjugacy classes, we have

$$\sum_{g^\sharp\in G^\sharp}|g^\sharp|=|G|.$$

We isolate the contribution of the conjugacy classes of size 1, i.e., of the center $Z(G)$ of G, and then reduce modulo p to get

$$\sum_{g^\sharp\in G^\sharp-Z(G)}|g^\sharp|\equiv-|Z(G)|\,(\mathrm{mod}\,p).$$

Thus either the center of G is not reduced to 1 (and we can take any non-trivial element $g\in Z(G)$, with $p\nmid|g^\sharp|=1$), or else one of the terms in the left-hand side, must be non-zero modulo p, and using such a g we can ensure all of (4.55). □

Remark 4.7.6 (Why A_5 is not solvable...). The first (in terms of order!) non-solvable group is the alternating group A_5 of order $60=2^2\cdot 3\cdot 5$ (one can show that all groups of order 30 – there are four up to isomorphism – are solvable.) It is instructive to see "how" the argument fails in that case. The character table of A_5 is computed, e.g., in [**13**, §3.1, Ex. 3.5], and we just list it here, subscripting the conjugacy classes with their sizes (the reader who has not seen it might think of finding natural linear actions corresponding to the representations displayed):

	1_1	$(12)(34)_{15}$	$(123)_{20}$	$(12345)_{12}$	$(13452)_{12}$
$\mathbf{1}$	1	1	1	1	1

	1_1	$(12)(34)_{15}$	$(123)_{20}$	$(12345)_{12}$	$(13452)_{12}$
ϱ_3	3	-1	0	$\frac{1+\sqrt{5}}{2}$	$\frac{1-\sqrt{5}}{2}$
ϱ_3'	3	-1	0	$\frac{1-\sqrt{5}}{2}$	$\frac{1+\sqrt{5}}{2}$
ϱ_4	4	0	1	-1	-1
ϱ_5	5	1	-1	0	0

Table 4.6. Character table of A_5

The pairs (g, ϱ) for which $\chi_\varrho(g) \neq 0$ and (4.55) holds are the following:

$$(p, q) = (2, 3), \qquad (g^\sharp, \varrho) = ((12)(34), \varrho_4),$$
$$(p, q) = (3, 2), \qquad (g^\sharp, \varrho) = ((123), \varrho_5),$$
$$(p, q) = (2, 5), \qquad (g^\sharp, \varrho) = ((12)(34), \varrho_3 \text{ or } \varrho_3'),$$
$$(p, q) = (5, 2), \qquad (g^\sharp, \varrho) = ((12345) \text{ or } (13452), \varrho_3 \text{ or } \varrho_3')$$
$$(p, q) = (3, 5), \qquad (g^\sharp, \varrho) = ((123), \varrho_4)$$
$$(p, q) = (5, 3), \qquad (g^\sharp, \varrho) = ((12345) \text{ or } (13452), \varrho_4).$$

As it should be, one sees that in all cases, the actual gcd of $|g^\sharp|$ and $\dim(\varrho)$ is different from 1.

We now come to the proof of Theorem 4.7.5. Here again, the basic ingredeitn is of independent interest, as it provides more subtle integrality properties of character values:

Proposition 4.7.7 (Divisibility). *Let G be a finite group and let*

$$a = \sum_{g \in G} \alpha(g) g \in \mathbf{C}[G]$$

be an element of the group ring with coefficients $\alpha(g)$ which are algebraic integers. *Moreover, assume $a \in Z(\mathbf{C}(G))$ is in the center of the group ring, or equivalently that α is a class function on G. Then a acts on irreducible representations by multiplication by scalars, and those are all algebraic integers.*

In particular, for any $g \in G$ and $\varrho \in \hat{G}$, we have

$$\dim(\varrho) \mid \chi_\varrho(g)|g^\sharp| \tag{4.58}$$

in the ring $\bar{\mathbf{Z}}$ of algebraic integers.

Proof. The action of a central element a on an irreducible representation ϱ is given by the scalar $\omega_\varrho(a)$ of Proposition 4.3.30, and so we must show that this is in $\bar{\mathbf{Z}}$ under the stated conditions.

Since, as a function of a, this scalar is a ring-homomorphism, and $\bar{\mathbf{Z}}$ is itself a ring (Proposition A.1.2), it is enough to prove the integrality of $\omega_\varrho(a)$ when a runs over a set of elements which span the subring $Z(\bar{\mathbf{Z}}[G])$. For instance, one can take the elements

$$a_c = \sum_{g \in c} g$$

where c runs over conjugacy classes in G. This means, in practice, that we may assume that the coefficients $\alpha(g)$ are in fact in $\mathbf{Z}$.

Now, under this condition, we consider the element $e_\varrho \in \mathbf{C}[G]$ giving the ϱ-isotypic projector. Using the left-multiplication action of G on the group ring, we have

$$ae_\varrho = \omega_\varrho(a)e_\varrho,$$

i.e., multiplication by a, as a map on $\mathbf{C}[G]$, has $\omega_\varrho(a)$ as an eigenvalue. But now we claim that this linear map

$$\Phi_a \begin{cases} \mathbf{C}[G] & \longrightarrow & \mathbf{C}[G] \\ x & \mapsto & ax \end{cases}$$

can be represented by an integral matrix in a suitable basis. In fact, the elements $x \in G$ form a basis in $\mathbf{C}[G]$ which does the job: we have

$$ax = \sum_{g\in G} \alpha(g)gx = \sum_{g\in G} \alpha(gx^{-1})g$$

where the relevant coefficients, namely the $\alpha(gx^{-1})$, are indeed integers.

We conclude that $\omega_\varrho(a)$ is a root of the characteristic polynomial $\det(X - \Phi_a)$ of Φ_a, and if we use the basis above, this is monic with integral coefficients, showing that $\omega_\varrho(a)$ is indeed an algebraic integer. (We are using here one part of the criterion in Proposition A.1.2.)

Now for the last part, we use the expression (4.26) for $\omega_\varrho(a)$, in the special case where $a = a_c$, which is

$$\omega_\varrho(a_c) = \frac{1}{\dim(\varrho)} \sum_{g\in c} \chi_\varrho(g) = \frac{|g^\sharp|\chi_\varrho(g)}{\dim(\varrho)},$$

and the fact that this is an algebraic integer is equivalent with the divisibility relation (4.58). □

Before using this to conclude, the reader is probably tempted to apply the general fact that $\omega_\varrho(a) \in \bar{\mathbf{Z}}$ to other elements a. In fact this gives the proof of an observation we already mentioned (see Remark 4.3.7 for instance):

Proposition 4.7.8 (Dimensions of irreducible representations divide the order). *If G is a finite group and $\varrho \in \hat{G}$ is an irreducible complex representation of G, the dimension of ϱ divides $|G|$.*

Proof. We are looking for a suitable $a \in \mathbf{C}[G]$ to apply the proposition; since

$$\omega_\varrho(a) = \frac{1}{\dim(\varrho)} \sum_{g\in G} \alpha(g)\chi_\varrho(g),$$

the most convenient would be to have $\alpha(g)$ such that the sum is equal to $|G|$. But there does exist such a choice: by the orthogonality relation, we can take $\alpha(g) = \overline{\chi_\varrho(g)}$ and then

$$\omega_\varrho(a) = \frac{1}{\dim(\varrho)} \sum_{g\in G} \alpha(g)\chi_\varrho(g) = \frac{|G|}{\dim(\varrho)}.$$

Since $\alpha(g) \in \bar{\mathbf{Z}}$, this is indeed an algebraic integer, by the proposition. Hence

$$\frac{|G|}{\dim(\varrho)} \in \bar{\mathbf{Z}} \cap \mathbf{Q} = \mathbf{Z},$$

which is the desired result. □

We can finally finish:

Proof of Theorem 4.7.5. With (4.58) in hand, what to do is quite clear: the dimension $\dim(\varrho)$ divides the product

$$\chi_\varrho(g)|g^\sharp|,$$

and it is assumed that it is coprime with the second factor $|g^\sharp|$ are coprime. So it must divide the first, i.e., the character value $\chi_\varrho(g)$ (this happens in the ring $\bar{\mathbf{Z}}$ of algebraic integers, always; we are using Proposition A.1.5.)

Such a relation, we claim, is in fact equivalent with the conclusion of the theorem. This would again be clear if $\chi_\varrho(g)$ were in $\mathbf{Z}$, since the bound

$$|\chi_\varrho(g)| \leqslant \dim(\varrho)$$

and the divisibility $\dim(\varrho) \mid \chi_\varrho(g) \in \mathbf{Z}$ lead to

$$\chi_\varrho(g) \in \{-\dim(\varrho), 0, \dim(\varrho)\},$$

and we know that $|\chi_\varrho(g)| = \dim(\varrho)$ is equivalent with $\varrho(g)$ being a scalar matrix, i.e., $g \in \ker \bar{\varrho}$ (Proposition 4.6.4).

To deal with the general case, we must be careful because if we have non-zero algebraic integers z_1, z_2 with

$$z_1 \mid z_2,$$

we can not always conclude that $|z_1| \leqslant |z_2|$ (e.g., take $z_1 = 1$ and $z_2 = -1 + \sqrt{2}$.) What we do is take the norm on both sides of the divisibility relation and obtain

$$\dim(\varrho)^r \mid N(\chi_\varrho(g))$$

where r is the number of conjugates of $\chi_\varrho(g)$ (this is Corollary A.1.8). This is now a divisibility relation among integers, and if $\chi_\varrho(g) \neq 0$, we deduce the inequality

$$\dim(\varrho)^r \leqslant |N(\chi_\varrho(g))|.$$

Now, each conjugate of $\chi_\varrho(g)$ is a sum of $\dim(\varrho)$ roots of unity,[16] hence is of modulus $\leqslant \dim(\varrho)$. This means that

$$|N(\chi_\varrho(g))| \leqslant \dim(\varrho)^r,$$

and by comparison we must have equality in all the terms of the product, in particular

$$|\chi_\varrho(g)| = \dim(\varrho),$$

which – as before – gives $g \in \ker \bar{\varrho}$. □

Remark 4.7.9. We used the divisibility relation (4.58) in the previous proof by assuming that $\dim(\varrho)$ is coprime with the factor $|g^\sharp|$ on the right-hand side. What happens if we assume instead that $\dim(\varrho)$ is coprime with the second factor, $\chi_\varrho(g)$? One gets the conclusion that $\dim(\varrho)$ divides the size of the conjugacy class of g. This is of some interest; in particular, if there exists some g with $\chi_\varrho(g) = \pm 1$ (or even $\chi_\varrho(g)$ a root of unity), we have

$$\dim(\varrho) \mid |g^\sharp|.$$

We can see this "concretely" in the Steinberg representations of $\mathrm{GL}_2(\mathbf{F}_p)$, of dimension p: the values at semisimple conjugacy classes are roots of unity, and indeed $\dim(\mathrm{St}) = p \mid p(p+1)$, $p(p-1)$, which are the sizes of the split (resp. non-split) semisimple classes. On the other hand, it is not clear if this "dual"statement has any interesting applications in group theory...

[16] In fact, it is also a character value $\chi_\varrho(x)$ for some $x \in G$, but checking this fact requires the Galois-theoretic interpretation of the conjugates, which we do not wish to assume.

Remark 4.7.10 (Characters have zeros). We come back to the following observation, which is certainly experimentally true for those groups for which we computed the character table:

Proposition 4.7.11 (Burnside). *Let G be a finite group, and $\varrho \in \hat{G}$ an irreducible representation of dimension at least* 2. *Then there exists some $g \in G$ such that $\chi_\varrho(g) = 0$.*

Proof. This is once again obvious if the character takes actual integer values in $\mathbf{Z}$: the orthonormality relation for ϱ gives

$$\frac{1}{|G|}\sum_{g\in G}|\chi_\varrho(g)|^2 = 1,$$

i.e., the mean-square average over G of the character of ϱ is 1. Hence either $|\chi_\varrho(g)|^2 = 1$ for all g, which can only happen when $\dim(\varrho) = 1$, or else some element g must have

$$|\chi_\varrho(g)| < 1.$$

If $\chi_\varrho(g) \in \mathbf{Z}$, this gives immediately $\chi_\varrho(g) = 0$. In the general case, we must again be careful, since there are many non-zero algebraic integers z with $|z| < 1$ (e.g., $-1 + \sqrt{2}$). However, one can partition G into subsets for which the sum of the character values is an actual integer. To be precise, we write G as the union of the equivalence classes for the relation defined by $x \sim y$ if and only if x and y generate the same (finite cyclic) subgroup of G. Hence

$$\sum_{g\in G}|\chi_\varrho(g)|^2 = \sum_{S\in G/\sim}\sum_{x\in S}|\chi_\varrho(x)|^2.$$

Each class S is the set of generators of some finite cyclic subgroup H of G. Applying (to H and the restriction of ϱ to H) the inequality (4.38) from Exercise 4.5.7, we deduce that

$$\sum_{x\in S}|\chi_\varrho(x)|^2 \geqslant |S|$$

unless some (in fact, all) character values $\chi_\varrho(x)$ are zero for $x \in S$. Summing over S, and comparing with the orthonormality relation, it follows that when χ_ϱ has no zero, there must be equality in each of these inequalities. But $S = \{1\}$ is one of the classes, and therefore $|\chi_\varrho(1)|^2 = 1$, which gives the desired result by contraposition. □

This fact is about the rows of the character table; is there another, "dual", property of the columns? If there is, it is not the existence of at least one zero entry in each column, except for those of central elements (for which the modulus of the character value is the dimension): although this property holds in a number of examples, for instance the groups $\mathrm{GL}_2(\mathbf{F}_p)$, we can see that it is false for the solvable groups B_p of Section 4.6.3: we see in Table 4.3 that for the non-diagonalizable elements $\begin{pmatrix} a & 1 \\ 0 & a \end{pmatrix}$, with conjugacy classes of size $p-1$, every character value is a root of unity.

4.7.3. Relations between roots of polynomials. Our last application is to a purely algebraic problem about polynomials: given a field k (arbitrary to begin with) and a non-zero irreducible polynomial $P \in k[X]$ of degree $d \geqslant 1$, the question is whether the roots

$$x_1, \ldots, x_d$$

of P (in some algebraic closure of k) satisfy any non-trivial linear, or multiplicative, relation? By this, we mean, do there exist coefficients $\alpha_i \in k$, not all zero, such that

$$\alpha_1 x_1 + \cdots + \alpha_d x_d = 0,$$

(linear relation) or integers $n_i \in \mathbf{Z}$, not all zero, such that

$$x_1^{n_1} \cdots x_d^{n_d} = 1.$$

For instance, since

$$x_1 + \cdots + x_d = a_{d-1}, \quad x_1 \cdots x_d = (-1)^d a_0,$$

for

$$P = X^d + a_{d-1}X^{d-1} + \cdots + a_1 X + a_0,$$

we have a non-trivial linear relation

$$x_1 + \cdots + x_d = 0,$$

whenever the coefficient of degree $d-1$ of P is zero, and a non-trivial multiplicative relation

$$x_1^2 \cdots x_d^2 = 1$$

whenever $a_0 = P(0) = \pm 1$.

A general method to investigate such questions was found by Girstmair (see [**15**] and the references there, for instance), based on representation theory. We present here the basic idea and the simplest results.

As in the (first) proof of Burnside's Irreducibility Criterion (Section 2.7.3), the basic idea is to define the set of all relations (linear or multiplicative) between the roots of P, and show that it is carries a natural representation of a certain finite group G. If we can decompose this representation in terms of irreducible representations, we will obtain a classification of the possible relations that can occur. As was the case for the Burnside criterion, this is often feasible because the relation space is a subrepresentation of a well-understood representation of G.

First, with notation as before for the roots of P, we denote by

$$R_a = \{(\alpha_i)_i \in k^d \mid \sum_{i=1}^{d} \alpha_i x_i = 0\},$$

$$R_m = \{(n_i)_i \in \mathbf{Z}^d \mid \prod_{i=1}^{d} x_i^{n_i} = 1\}$$

the spaces of linear or multiplicative relations between the roots; we see immediately that R_a is a k-vector subspace of k^d, while R_m is a subgroup of the abelian group $\mathbf{Z}^d$, so that it is a free abelian group of rank at most d.

The group G that acts naturally on these spaces is the *Galois group* of the polynomial P, which means the Galois group of its splitting field

$$k_P = k(x_1, \ldots, x_d).$$

To ensure that this Galois group is well-defined, we must assume that P is separable, for instance that k has characteristic zero ($k = \mathbf{Q}$ will do). The elements of G are therefore field automorphisms

$$\sigma : k_P \longrightarrow k_P.$$

By acting on a relation (linear or multiplicative) using G, we see that the Galois group acts indeed on R_a and R_m. More precisely, recall that $\sigma \in G$ permutes the roots (x_i), so that there exists a group homomorphism

$$\begin{cases} G & \longrightarrow & \mathfrak{S}_d \\ \sigma & \mapsto & \hat{\sigma} \end{cases}$$

characterized by

$$\sigma(x_i) = x_{\hat{\sigma}(i)}$$

for all roots of P. This homomorphism is faithful since the roots of P generate the splitting field k_P.

If $\boldsymbol{\alpha} = (\alpha_i)$ is in R_a, acting by σ on the relation

$$\alpha_1 x_1 + \cdots + \alpha_d x_d = 0,$$

we get

$$0 = \sigma(\alpha_1 x_1 + \cdots + \alpha_d x_d) = \alpha_1 x_{\hat{\sigma}(1)} + \cdots + \alpha_d x_{\hat{\sigma}(d)}$$

(since σ is the identity on k) or in other words the vector

$$\sigma \cdot \boldsymbol{\alpha} = (\alpha_{\hat{\sigma}^{-1}(i)})_{1 \leqslant i \leqslant d}$$

is also in R_a. But note that we can define

$$\sigma \cdot \boldsymbol{\alpha} = (\alpha_{\hat{\sigma}^{-1}(i)})_{1 \leqslant i \leqslant d}$$

for *arbitrary* $\alpha \in k^d$, and $\sigma \in G$; this is in fact simply the *permutation k-representation* of G on k^d constructed from the action of G on the set $\{x_i\}$ of roots of G (see Section 2.6.2), and hence we see that R_a is a subrepresentation of this permutation representation, which we denote π_k.

Similarly, we can act with G on multiplicative relations. However, since R_m is only an abelian group, it is only the $\mathbf{Q}$-vector space $R_m \otimes \mathbf{Q}$ that one can see as a subrepresentation of the (same!) permutation representation $\pi_{\mathbf{Q}}$ of G over $\mathbf{Q}$ (one may also use any field containing $\mathbf{Q}$).

If we can decompose π_k, we can hope to see which subrepresentations can arise as relation spaces R_a, and similarly for $\pi_{\mathbf{Q}}$ and R_m. The simplest case is the following:

Proposition 4.7.12. *Let k be a field of characteristic zero, $P \in k[X]$ an irreducible polynomial of degree $d \geqslant 2$ with Galois group isomorphic to the full symmetric group $\mathfrak{S}_d$.*

(1) *Either $R_a = 0$, i.e., there are no non-trivial linear relations between the roots of P, or R_a is one-dimensional and is spanned by the element $e_0 = (1, \ldots, 1)$ corresponding to the relation*

$$x_1 + \cdots + x_d = 0.$$

This second case may always happen, for a given field k, if there exists a polynomial with Galois group $\mathfrak{S}_d$.

(2) *Either $R_m = 0$, i.e., there are non non-trivial multiplicative relations between the roots of P, or R_a is a free $\mathbf{Z}$-module of rank 1 generated by ne_0 for some $n \geqslant 1$, or the splitting field of P is contained in the splitting field of a Kummer polynomial $X^n - b$. The first case can happen for any field k for which there exists a polynomial with Galois group $\mathfrak{S}_d$, and the second and third cases are possible for $k = \mathbf{Q}$ for $n = 2$, $n = 3$.*

Proof. As we have already observed, the space k^d of π_k decomposes as a direct sum of subrepresentations

$$k^d = ke_0 \oplus V$$

where

$$V = \{v = (v_i) \in k^d \mid \sum_i v_i = 0\}.$$

When $G \simeq \mathfrak{S}_d$, although k is not algebraically closed (otherwise an irreducible P of degree $\geqslant 2$ would not exist!), these are irreducible subrepresentations. Indeed, we must only check this for V, and we immediately see that if V could be decomposed into two or more subrepresentations (recall that Maschke's theorem does apply for any field of characteristic 0), then the same would be true for the representation of G on $V \otimes \bar{k}$, which contradicts the fact that it is irreducible (though we have only directly proved this for $k \subset \mathbf{C}$, see (4.19)).

Because ke_0 and V are non-isomorphic as representations of $G \simeq \mathfrak{S}_d$ (even for $d = 2$, where V is also one-dimensional), the only possibilities for R_a are therefore

$$R_a = 0, \text{ or } ke_0, \text{ or } V, \text{ or } k^d$$

(this is the uniqueness of isotypic subspaces, see the second part of Proposition 2.7.7.) The cases $R_a = 0$ and $R_a = ke_0$ are precisely the two possibilities of the statement we try to prove, and we can now check that the others can not occur. For this, it is enough to show that $R_a \supset V$ is impossible. But V is spanned by the vectors

$$f_2 = (1, -1, 0, \ldots, 0), \quad \ldots, \quad f_d = (1, 0, \ldots, 0, -1). \tag{4.59}$$

Even if only the first were to be in R_a, this would translate into $x_1 = x_2$, which is impossible.

There only remains to prove the existence part: provided *some* polynomial in $k[X]$ with Galois group $\mathfrak{S}_d$ exists,[17] one can find instances where both cases $R_a = 0$ or $R_a = ke_0$ occur. This is easy: if we fix such a polynomial

$$P = X^d + a_{d-1}X^{d-1} + \cdots + a_1 X + a_0 \in k[X]$$

with Galois group $\mathfrak{S}_d$, the polynomials $P(X + a)$, for any $a \in k$, also have the same Galois group (the splitting field has not changed!), and the sum of the roots $y_i = x_i - a$ is

$$\sum_{i=1}^{d} x_i - da,$$

which takes all values in k when a varies; when it is zero, the polynomial $P(X + a)$ satisfies $R_a = ke_0$.

We now deal with the multiplicative case; the argument is similar, but some of the excluded cases become possible. First, since $R_m \otimes \mathbf{Q}$ is a subrepresentation of $\pi_{\mathbf{Q}}$, we see again that there are the same four possibilities for the subspace $R_m \otimes \mathbf{Q}$. Of course, if it is zero, we have $R_m = 0$ (because R_m is a free abelian group); if $R_m \otimes \mathbf{Q} = \mathbf{Q}e_0$, on the other hand, we can only conclude that $R_m = n\mathbf{Z}e_0$ for some integer $n \geqslant 1$ (examples below show that one may have $n \neq 1$.)

Continuing with the other possibilities, we have $R_m \otimes \mathbf{Q} \supset V$ if and only, for some $n \geqslant 1$, the vectors $nf_2, \ldots, nf_d$ are in R_m, where f_i is defined in (4.59). This means that we have

$$\Bigl(\frac{x_1}{x_2}\Bigr)^n = \cdots = \Bigl(\frac{x_1}{x_d}\Bigr)^n = 1,$$

and from this we deduce that

$$\sigma(x_1^n) = x_{\hat{\sigma}(1)}^n = x_1^n,$$

for all $\sigma \in G$. By Galois theory, this translates to $x_1^n \in k$. Therefore x_1 is a root of a *Kummer polynomial* $X^n - b \in k[X]$, which is the last possible conclusion we claimed. Note that b could be a root of unity (belonging to k): this is a special case, which corresponds to $R_m \otimes \mathbf{Q} = \mathbf{Q}^d$ (instead of V), since each x_i is then a root of unity.

In terms of existence of polynomials with these types of multiplicative relations, we first note that $R_m = 0$ is always possible if there exists at least *one* irreducible polynomial $P \in k[X]$ with Galois group $\mathfrak{S}_d$. Indeed, we have $P(0) \neq 0$, and as before, we may replace

[17] This may not be the case, or only for some d (in the case of $k = \mathbf{R}$, only $d = 2$ is possible.)

P with $Q = P(aX)$ for $a \in k$, without changing the Galois group; the roots of Q are $y_i = a^{-1}x_i$, and

$$y_1 \cdots y_d = \frac{x_1 \cdots x_d}{a^d}.$$

Then, if we pick $a \in k^\times$ so that this is expression is not a root of unity, we obtain a polynomial Q with $R_m = 0$ (such an $a \neq 0$ exists: otherwise, taking $a = 1$ would show that $x_1 \cdots x_d$ is itself a root of unity, and then it would follow that *any* $a \in k^\times$ is a root of unity, which is absurd since $\mathbf{Q} \subset k$).

For the case of $R_m \otimes \mathbf{Q} = \mathbf{Q}e_0$, we will just give examples for $k = \mathbf{Q}$: it is known (see, e.g., [**37**, p. 42]) that the polynomial

$$P = X^d - X - 1$$

has Galois group $\mathfrak{S}_d$ for $d \geqslant 2$; since the product of its roots is $(-1)^d P(0) = (-1)^{d+1}$, it satisfies the relation

$$\prod_i x_i^2 = 1,$$

so $R_m \otimes \mathbf{Q} = \mathbf{Q}e_0$, and in fact $R_m = n\mathbf{Z}e_0$ with $n = 1$ if d is odd, and $n = 2$ if d is even.

For the Kummer cases, we take $k = \mathbf{Q}$ for simplicity (it will be clear that many fields will do). For $n = 2$, any quadratic polynomial $X^2 - b$ with b not a square of a rational number has Galois group $\mathfrak{S}_2$; if $b = -1$, noting $x_1 = i$, $x_2 = -i$, we have

$$R_m = \{(n_1, n_2) \in \mathbf{Z}^2 \mid n_1 + 2n_2 \equiv 0 \pmod 4\},$$

which has rank 2 (so $R_m \otimes \mathbf{Q} = \mathbf{Q}^2$), and if $b \neq -1$, we have

$$R_m = \{(n_1, n_2) \in \mathbf{Z}^2 \mid n_i \equiv 0 \pmod 2),\ n_1 + n_2 = 0\},$$

with $R_m \otimes \mathbf{Q} = \mathbf{Q}e_0$. For $n = 3$, any Kummer equation $X^3 - b = 0$, with b not a perfect cube, will have splitting field with Galois group $\mathfrak{S}_3$, and a quick computation with the roots $\sqrt[3]{b}$, $j\sqrt[3]{b}$, $j^2\sqrt[3]{b}$, where j is a primitive cube root of unity in $\mathbf{C}$, leads to

$$R_m = \{(n_1, n_2, n_3) \in \mathbf{Z}^3 \mid n_1 + n_2 + n_3 = 0, \quad n_i \equiv n_j \pmod 3 \text{ for all } i, j\},$$

so that again $R_m \otimes \mathbf{Q} = \mathbf{Q}e_0$. □

Exercise 4.7.13 (Palindromic polynomials). We consider in this exercise the case where d is even and the Galois group of P is the group W_d defined as the subgroup of $\mathfrak{S}_d$ that respects a partition of $\{1, \ldots, d\}$ into $d/2$ pairs. More precisely, let X be a finite set of cardinality d, and let $i : X \to X$ be an involution on X (i.e., $i \circ i = \mathrm{Id}_X$) with no fixed points, for instance $X = \{1, \ldots, d\}$ and $i(x) = d + 1 - x$. The d pairs $\{x, i(x)\}$ partition X, and

$$W_d = \{\sigma \in \mathfrak{S}_d \mid \sigma(i(x)) = i(\sigma(x)) \text{ for all } x \in X\}$$

which means concretely that an element of W_d permutes the pairs $\{x, i(x)\}$, and may (or not) switch x and $i(x)$. This group is sometimes called the group of *signed permutations* of $\{1, \ldots, d\}$.

(1) Show that W_d is of order $2^d d!$ for $d \geqslant 2$ even, and that there is an exact sequence

$$1 \longrightarrow (\mathbf{Z}/2\mathbf{Z})^d \longrightarrow W_d \longrightarrow \mathfrak{S}_d \longrightarrow 1.$$

Find a faithful representation of W_d in $\mathrm{GL}_d(\mathbf{C})$ where the matrix representation has values in $\mathrm{GL}_d(\mathbf{Z})$. (See Example 7.1.2.)

(2) Let k be a field of characteristic 0, $P \in k[X]$ an irreducible polynomial of degree $d = 2n$ even, $d \geqslant 2$, of the form

$$P = X^d + a_{d-1}X^{d-1} + \cdots + a_{n+1}X^n + a_{n+1}X^{n-1} + \cdots + a_{d-1}X + 1,$$

i.e., with the same coefficients for X^j and X^{d+1-j} for all j (such polynomials are called *palindromic*, or *self-dual*). Show that the Galois group of P can be identified with a subgroup of W_d. [Hint: If x is a root of P, then $1/x$ is also one, and $1/x \neq x$, so that one can take $X = \{\text{roots of P}\}$ and $i(x) = 1/x$.]

(3) Assume that the Galois group of P is equal to W_d. Show that the permutation representation π_k of W_d associated to the action of W_d on X splits as a direct sum

$$E_0 \oplus E_1 \oplus E_2$$

where $E_0 = ke_0$ and, for a suitable numbering of the roots, we have

$$E_1 = \Big\{(\alpha_i) \mid \alpha_{d+1-i} - \alpha_i = 0,\ 1 \leqslant i \leqslant d, \quad \sum \alpha_i = 0\Big\},$$
$$E_2 = \Big\{(\alpha_i) \mid \alpha_{d+1-i} + \alpha_i = 0,\ 1 \leqslant i \leqslant d\Big\},$$

and the three spaces are irreducible. [Hint: You may assume $k \subset \mathbf{C}$; compute the orbits of W_d on X, and deduce the value of the squared norm of the character of π_k.]

(4) Show that the only possible spaces of linear relation between roots of a polynomial P as above are $R_a = 0$ and $R_a = ke_0$.

It is known that "many" palindromic polynomials with Galois groups W_d exist; for more information and some applications, the reader may look at [**25**, §2].

Remark 4.7.14. Both Proposition 4.7.12 and this exercise are in the direction of showing that linear or multiplicative relations are rare in some cases. However, for some Galois groups, interesting things can happen. For instance, Girstmair showed that there exists a group G of order 72 which can arise as the Galois group (for $k = \mathbf{Q}$) of some polynomial P of degree 9 for which the roots, suitably numbered, satisfy

$$4x_1 + x_2 + x_3 + x_4 + x_5 - 2(x_6 + x_7 + x_8 + x_9) = 0.$$

Another example is the group G usually denoted $W(E_8)$, the "Weyl group of E_8", which can be defined as the group with 8 generators

$$w_1, \ldots, w_8$$

which are subject to the relations

$$w_i^2 = 1 \qquad (w_iw_j)^{m(i,j)} = 1, \qquad 1 \leqslant i < j \leqslant 8,$$

where

$$m(i,j) = 3 \text{ if } (i,j) \in \{(1,3),\ (3,4),\ (2,4),\ (4,5),\ (5,6),\ (6,7),\ (7,8)\},$$

and $m(i,j) = 2$ otherwise. (This definition is given here only in order to be definite; of course, this presentation is justified by the many other definitions and properties of this group, which is a special case of *Coxeter groups*.) The group $W(E_8)$ has order

$$W(E_8) = 696,729,600 = 2^{14} \cdot 3^5 \cdot 5^2 \cdot 7,$$

and one can construct irreducible polynomials $P \in \mathbf{Q}[X]$, of degree 240, with Galois group $W(E_8)$, such that

$$\dim(R_m \otimes \mathbf{Q}) = 232,$$

or in other words: there are 8 roots of P, out of 240, such that all others are in the multiplicative group generated by those (see [**2**, §5] or [**21**, Rem. 2.4]).

4.8. Further topics

We finish this chapter with a short discussion of some further topics concerning representations of finite groups. These – and their developments – are of great interest and importance in some applications, and although we only consider basic facts, we will give references where more details can be found. One last important notion, the *Frobenius-Schur indicator* of an irreducible representation, will be considered in Section 6.2, because it makes sense and is treated exactly the same way for all compact groups.

4.8.1. More on induction. We have used induction quite often, either in a general way (typically to exploit Frobenius reciprocity) or to construct specific representations of concrete groups. In the second role, in particular, we see that it is useful to understand intertwiners between two induced representations. In particular in Section 4.6.4, we computed the dimension of such spaces "by hand", as inner products of induced characters. The answers are rather clean, as (4.49), and it should not be a surprise to see that there is a general approach to these computations.

Proposition 4.8.1 (Intertwiners between induced representations). *Let G be a finite group, and let H_1, H_2 be subgroups of G, and ϱ_1, ϱ_2 complex finite-dimensional representations of H_1 and H_2, acting on the vector spaces E_1, and E_2,respectively. There is an isomorphism*

$$\mathrm{Hom}_G(\mathrm{Ind}_{H_1}^G \varrho_1, \mathrm{Ind}_{H_2}^G(\varrho_2)) \simeq I_{\varrho_1,\varrho_2}$$

where

$$I_{\varrho_1,\varrho_2} = \{\alpha : G \to \mathrm{Hom}_{\mathbf{C}}(E_1,E_2) \mid \alpha(h_1xh_2) = \varrho_2(h_2)^{-1}\circ\alpha(x)\circ\varrho_1(h_1)^{-1},\\ \textit{for all } h_1 \in H_1,\ x \in G,\ h_2 \in H_2\}. \tag{4.60}$$

Proof. We start naturally by applying Frobenius reciprocity to "remove" one induced representation: we have the isomorphism

$$\mathrm{Hom}_G(\mathrm{Ind}_{H_1}^G \varrho_1, \mathrm{Ind}_{H_2}^G(\varrho_2)) \simeq \mathrm{Hom}_{H_2}(\mathrm{Res}_{H_2}^G \mathrm{Ind}_{H_1}^G \varrho_1, \varrho_2).$$

The idea now is to find first a convenient model for the space

$$\mathrm{Hom}(\mathrm{Res}_{H_2}^G \mathrm{Ind}_{H_1}^G \varrho_1, \varrho_2),$$

and then to isolate inside the H_2-intertwiners. Let F_1 denote the space on which $\mathrm{Ind}_{H_1}^G \varrho_1$ acts, as well as its restriction to H_2. By construction, F_1 is a subspace of the space V_1 of all functions from G to E_2, and hence we can write any linear map T from F_1 to E_2 in the form

$$T\varphi = T_\alpha\varphi = \sum_{x\in G}\alpha(x)(\varphi(x))$$

for some $\alpha(x) \in \mathrm{Hom}(E_1,E_2)$ (this amounts to using the basis of characteristic functions of single points to compute linear maps with values in E_2 which are defined on the whole of V_1). We claim that this gives an isomorphism

$$T : \begin{cases} I & \longrightarrow & \mathrm{Hom}(F_1,E_2) \\ \alpha & \mapsto & T_\alpha \end{cases}$$

where

$$I = \{\alpha : G \longrightarrow \mathrm{Hom}(E_1,E_2) \mid \alpha(h_1x) = \varrho_1(h_1)^{-1}\alpha(x)\}$$

(in general, since F_1 is a subspace of V_1, $\mathrm{Hom}(F_1,E_2)$ would be a quotient of $\mathrm{Hom}(V_1,E_2)$, and what we are doing is find a good representative subspace for it in $\mathrm{Hom}(V_1,E_2)$).

Indeed, the computation

$$\begin{aligned}T_\alpha(\varphi) &= \sum_{y\in H_1\backslash G}\sum_{h_1\in H_1}\alpha(h_1y)(\varphi(h_1y))\\ &= \sum_{y\in H_1\backslash G}\sum_{h_1\in H_1}\alpha(y)\circ\varrho_1(h_1)^{-1}(\varrho_1(h_1)\varphi(y))\\ &= |H_1|\sum_{y\in H_1\backslash G}\alpha(y)(\varphi(y))\end{aligned}$$

quickly shows that T is injective (since one can prescribe the values of an element $\varphi\in F_1$ arbitrarily on each coset in $H_1\backslash G$). But $\dim I = [G:H](\dim E_1)(\dim E_2)$ (by an argument similar to the computation of the dimension of an induced representation), and this is also the dimension of $\mathrm{Hom}(F_1,E_2)$ (by Proposition 2.3.8), so T is indeed an isomorphism.

Now we can easily answer the question: given $\alpha\in I$, when is $T_\alpha\in\mathrm{Hom}(F_1,E_2)$ an H_2-intertwiner? We have

$$T_\alpha(h_2\cdot\varphi) = \sum_{x\in G}\alpha(x)((h_2\cdot\varphi)(x)) = \sum_{x\in G}\alpha(x)(\varphi(xh_2))$$

and we want this to be equal to

$$\varrho_2(h_2)T_\alpha(\varphi) = \sum_{x\in G}(\varrho_2(h_2)\circ\alpha(x))(\varphi(x))$$

for all $h_2\in H_2$, $\varphi\in F_1$. We fix h_2; by change of variable, the first expression is

$$T_\alpha(h_2\cdot\varphi) = \sum_{x\in G}\alpha(xh_2^{-1})(\varphi(x)) = T_\beta(\varphi)$$

for $\beta(x)=\alpha(xh_2^{-1})$. The second is $T_{\varrho_2(h_2)\alpha}$, and since β and $\varrho_2(h_2)\alpha$ are both still elements of I, the injectivity of T shows that the equality

$$T_\alpha(h_2\cdot\varphi) = \varrho_2(h_2)T_\alpha(\varphi),\qquad \varphi\in F_1,$$

is equivalent to

$$\alpha(xh_2^{-1}) = \varrho_2(h_2)\alpha(x)$$

for all $x\in G$. Replacing h_2 by h_2^{-1}, and combining these for all $h_2\in H_2$, we find that I_{ϱ_1,ϱ_2} defined in (4.60) is the subspace of I isomorphic, under T, to $\mathrm{Hom}_{H_2}(F_1,E_2)$. □

If, as in the case of $\mathrm{GL}_2(\mathbf{F}_p)$ in Section 4.6.4, we induce one-dimensional characters, we get a very general irreducibility criterion:

Corollary 4.8.2 (Irreducibility of induced representations). *Let G be a finite group, H a subgroup of G and χ a one-dimensional complex representation of H. The induced representation $\varrho=\mathrm{Ind}_H^G\chi$ is irreducible if and only if the one-dimensional representations χ_s of $H_s=H\cap sHs^{-1}$ defined by*

$$\chi_s(h)=\chi(s^{-1}hs)$$

are distinct from $\mathrm{Res}_{H_s}^G\chi$ as s runs over the complement of H in G.

Proof. By the irreducibility criterion (Corollary 4.3.14, which is also the converse of Schur's Lemma), we need to determine when the space $I_{\chi,\chi}$ of intertwiners of $\varrho=\mathrm{Ind}_H^G(\chi)$ with itself is one-dimensional. We apply Proposition 4.8.1 to compute this space; if we note that the space $\mathrm{Hom}(E_1, E_2)$ can be identified with $\mathbf{C}$ when $E_1 = E_2$ is one-dimensional, we see that $I_{\chi,\chi}$ is isomorphic to the space I of functions $\alpha : G \longrightarrow \mathbf{C}$ such that

$$\alpha(h_1 x h_2) = \varrho(h_1 h_2)^{-1}\alpha(x)$$

for all $x \in G$ and h_1, $h_2 \in H$. These conditions seem similar to those defining an induced representation, and this would suggest at first that the dimension of I is $H\backslash G/H = |S|$, but there is a subtlety: a representation $x = h_1 s h_2$ with $h_i \in H$ need *not* be unique, which creates additional relations to be satisfied. In consequence, the dimension can be smaller than this guess.

The one-dimensional subspace of I corresponding to $\mathbf{C}\mathrm{Id}$ in $\mathrm{End}_G(\varrho)$ is spanned by the function α_0 such that

$$\alpha_0(x) = \begin{cases} \chi(x)^{-1} & \text{if } x \in H, \\ 0 & \text{otherwise,} \end{cases}$$

(as one can easily check using the explicit form of the Frobenius reciprocity isomorphism).

We now determine the condition under which I is actually spanned by this special function α_0. If we denote by S a set of representations for the double cosets $HsH \subset G$ (taking $s = 1$ for the double coset $H \cdot H = H$), we see first of all, from the relations defining I, that the map

$$\begin{cases} I & \longrightarrow & \mathbf{C}^S \\ \alpha & \mapsto & (\alpha(s))_{s\in S} \end{cases}$$

is *injective*. Similarly, we get

$$\alpha(hs) = \chi(h)^{-1}\alpha(s) = \alpha(sh) \tag{4.61}$$

for all $h \in H$ and $s \in S$.

Now fix $s \in S$. We claim that either $\alpha(s) = 0$ or $\chi_s = \chi$ on $H_s = H \cap sHs^{-1}$ (note in passing that χ_s is indeed a well-defined representation of this subgroup, since $s^{-1}H_s s \subset H$). Indeed, for $x \in H_s = H \cap sHs^{-1}$, we have

$$\alpha(xs) = \alpha(s(s^{-1}xs)) = \alpha((s^{-1}xs)s)$$

by (4.61), applied with $h = s^{-1}xs \in H$, and this gives

$$\chi(x)^{-1}\alpha = \chi(s^{-1}xs)^{-1}\alpha,$$

which is valid for all $x \in H_s$, hence the claim.

Consequently, if no χ_s coincides with χ on H_s when $s \notin H$ (corresponding to the double cosets which are distinct from H), any $\alpha \in I$ must vanish on the non-trivial double cosets, and hence be a multiple of α_0. This proves the sufficiency part of the irreducibility criterion, and we leave the necessity to the reader... □

Exercise 4.8.3 (Principal series). Let $n \geqslant 2$ be an integer and let p be a prime number. Let $G = \mathrm{GL}_n(\mathbf{F}_p)$ and define $B \subset G$ to be the subgroup of all upper-triangular matrices. Moreover, let $W \subset G$ be the subgroup of permutation matrices.

(1) Show that

$$G = \bigcup_{w\in W} BwB,$$

and that this is a disjoint union. [Hint: Think of using Gaussian elimination to triangulate a matrix in some basis.]

(2) Let $\chi : B \longrightarrow \mathbf{C}^\times$ be the one-dimensional representation given by

$$\chi(g) = \chi_1(g_{1,1})\chi_2(g_{2,2})\cdots\chi_n(g_{n,n})$$

where χ_i, $1 \leqslant i \leqslant n$, are characters of $\mathbf{F}_p^\times$, and let $\varrho = \operatorname{Ind}_B^G \chi$. Show that ϱ is irreducible whenever all characters χ_i are distinct.

(3) For $n = 2$, show without using characters that the induced representations $\pi(\chi_1, \chi_2)$ and $\pi(\chi_2, \chi_1)$ (with $\chi_1 \neq \chi_2$) are isomorphic, and write down a formula for an isomorphism. [Hint: Follow the construction in Proposition 4.8.1 and the Frobenius reciprocity isomorphism.]

The irreducible representations of $\mathrm{GL}_n(\mathbf{F}_p)$ constructed in this exercise are called the *principal series*; for $n = 2$, they are exactly those whose irreducibility was proved in Section 4.6.4 using their characters.

Note that there is no principle series unless there are at least n distinct characters of $\mathbf{F}_p^\times$, i.e., unless $p - 1 \geqslant n$. This suggests – and this is indeed the case! – that the character tables of $\mathrm{GL}_n(\mathbf{F}_p)$, when p is fixed and n varies, behave rather differently from those of $\mathrm{GL}_n(\mathbf{F}_p)$ when n is fixed and p varies.

Exercise 4.8.4. We explain here a different approach to the corollary. Let G be a finite group, H_1 and H_2 subgroups of G, χ a one-dimensional complex representation of H_1.

(1) Show that

$$\operatorname{Res}_{H_2}^G \operatorname{Ind}_{H_1}^G \chi \simeq \bigoplus_{s \in S} \operatorname{Ind}_{H_{2,s}}^{H_2} \chi_s$$

where S is a set of representatives of the double cosets $H_2 g H_1$, $H_{2,s} = H_2 \cap s^{-1} H_1 s$ and χ_s is the one-dimensional character of $H_{2,s}$ given by

$$\chi_s(x) = \chi(sxs^{-1})$$

[Hint: This can be done with character theory.]

(2) Prove the corollary, and recover the irreducibility result (2) of the previous exercise, using (1).

Our second topic concerning induction takes up the following question: we have seen, in concrete examples, that many irreducible representations of certain finite groups arise as induced representations from subgroups, and indeed from one-dimensional characters of abelian subgroups. How general is this property? It turns out that, provided some leeway is allowed in the statement, one can in fact recover *all* irreducible representations using induced representations of cyclic subgroups. This is the content of the following theorem of Artin, which is an excellent illustration of the usefulness of the character ring $R(G) = R_{\mathbf{C}}(G)$ of virtual characters of G introduced in Definition 2.7.40.

Theorem 4.8.5 (Artin). *Let G be a finite group, and let $\varrho \in \hat{G}$ be an irreducible representation of G. There exists a decomposition*

$$\varrho = \sum_i \alpha_i \operatorname{Ind}_{H_i}^G \chi_i \tag{4.62}$$

in $R(G) \otimes \mathbf{Q}$, where $\alpha_i \in \mathbf{Q}$, $H_i \subset G$ is a cyclic subgroup of G and χ_i is an irreducible character of H_i, and where we identify representations of G with their image in $R(G)$.

Concretely, recall that the meaning of (4.62) is the following: there exist $m \geqslant 1$, finitely many non-negative integers $n_i \geqslant 0$ and $m_j \geqslant 0$, corresponding cyclic subgroups H_i and H_j, and characters χ_i and χ_j, such that we have an isomorphism of actual representations

$$m\varrho \oplus \bigoplus_i n_i \operatorname{Ind}_{H_i}^G \chi_i \simeq \bigoplus_j m_j \operatorname{Ind}_{H_j}^G \chi_j.$$

(precisely, m is a common denominator of all α_i in (4.62), $n_i = -m\alpha_i$ if $\alpha_i < 0$, while $m_j = m\alpha_j$ if $\alpha_j \geqslant 0$).

As we will see with examples, this statement can not, in general, be improved to express ϱ without considering virtual characters, i.e., with all α_i non-negative.

Proof. There is a surprisingly easy proof: the $\mathbf{Q}$-vector space $R(G) \otimes \mathbf{Q}$ has a basis corresponding to the irreducible representations of G, and inherits a non-degenerate symmetric bilinear form $\langle \cdot, \cdot \rangle_G$ for which those characters form an orthonormal basis. By duality of vector spaces, a subspace $V \subset R(G) \otimes \mathbf{Q}$ is equal to the whole space if and only if its orthogonal $V^{\perp}$ is zero. In particular, if V is generated by certain elements (χ_i) in $R(G) \otimes \mathbf{Q}$, we $V = R(G) \otimes \mathbf{Q}$ if and only if no $\xi \in R(G) \otimes \mathbf{Q}$ is orthogonal to all χ_i.

We apply this now to the family χ_i of all representations induced from irreducible representations of cyclic subgroups of G. Suppose $\xi \in R(G) \otimes \mathbf{Q}$ is orthogonal to all such χ_i. Then, using the Frobenius reciprocity formula (which holds in $R(G) \otimes \mathbf{Q}$ by "linearity" of induction and restriction with respect to direct sums), we get

$$\langle \mathrm{Ind}_H^G \chi, \xi \rangle_G = \langle \chi, \mathrm{Res}_H^G \xi \rangle_H$$

for any cyclic subgroup H and $\chi \in \hat{H}$. Varying χ for a fixed H, it follows that $\mathrm{Res}_H^G \xi$ is zero for all H. If we identiy ξ with its virtual character in $C(G)$, this means that ξ vanishes on all cyclic subgroups of G. But since any element $x \in G$ belongs to at least one cyclic subgroup, this gives $\xi = 0$. □

4.8.2. Representations over other fields.

CHAPTER 5

Abstract representation theory of compact groups

In this chapter, we consider the representation theory of compact topological groups. Our goal is to present the basic facts from the general theory, which is due to Peter-Weyl, and to do so by highlighting the close parallel with the case of finite groups (which is a special case, in fact, if we consider a finite group as a compact group with the discrete topology). This requires some care in the analytic set up, but the reader should appreciate how the work in getting the right definition of continuity, and of the regular representation (for instance) are justified when, in the end, the character formalism and the decomposition of the regular representation look formally very much the same as they do for finite groups.

5.1. An example: the circle group

We begin with an example, where it turns out that the basic facts are already well-known from the theory of Fourier series. This corresponds to what is probably the simplest infinite compact group, the unit circle

$$G = \{z \in \mathbf{C}^\times \mid |z| = 1\},$$

with its topology as a subset of $\mathbf{C}$. This groups is often best understood in the equivalent representation as the quotient $\mathbf{R}/\mathbf{Z}$, or $\mathbf{R}/2\pi\mathbf{Z}$, with the quotient topology, where the isomorphism is given by

$$\begin{cases} \mathbf{R}/\mathbf{Z} & \longrightarrow & G \\ x & \mapsto & e^{2i\pi x} \end{cases}$$

(since it is important to view G as a topological group, one should note that this is a *homeomorphism.*)[1]

We know already an infinite family of one-dimensional (in particular, irreducible) unitary representation of G, namely the characters

$$\chi_m : \begin{cases} \mathbf{R}/\mathbf{Z} & \longrightarrow & \mathbf{C}^\times \\ t & \mapsto & e^{2i\pi mt} \end{cases}$$

for $m \in \mathbf{Z}$. If we think of the space of functions on G, we see that linear combinations of these characters are simply trigonometric polynomials. This is a rather special subspace of functions, but it is dense in many important function spaces, including the space of continuous functions, by the Weierstrass approximation theorem.

The problem of expressing an "arbitrary" function in a series involving the χ_m is a basic problem of Fourier analysis, and is one of the most classical (and beautiful) problems of analysis. The Fourier series of an integrable function φ is the series

$$\sum_{m\in\mathbf{Z}} \alpha(m)e^{2i\pi mx} = \sum_{m\in\mathbf{Z}} \alpha(m)\chi_m(x)$$

[1] It is useful here to remember that a continuous bijection between compact topological spaces is necessarily a homeomorphism, i.e., the inverse is also automatically continuous.

where

$$\alpha(m) = \int_{\mathbf{R}/\mathbf{Z}} f(t) e^{-2i\pi mt} dt$$

are the Fourier coefficients. Many results are known about the convergence of Fourier series towards φ, many of which reveal of great subtlety of behavior. For instance, it was shown by Kolmogorov that there exist integrable functions $\varphi \in L^1(\mathbf{R}/\mathbf{Z})$ such that the Fourier series above diverges *for all* $x \in \mathbf{R}/\mathbf{Z}$. (For the classical theory of Fourier series, one can look at Zygmund's treatise [**45**].)

However, it is also classical that a very good theory emerges when considering the square-integrable functions: if $\varphi \in L^2(G)$, the Fourier series converges in L^2-norm, i.e.,

$$\Bigl\| \varphi - \sum_{|m| \leqslant M} \alpha(m) \chi_m \Bigr\|_{L^2} \longrightarrow 0$$

as $M \longrightarrow +\infty$, and in particular the Parseval formula

$$\int_{\mathbf{R}/\mathbf{Z}} |\varphi(x)|^2 dx = \sum_{m \in \mathbf{Z}} |\alpha(m)|^2 = \sum_{m \in \mathbf{Z}} |\langle \varphi, \chi_m \rangle|^2$$

holds.

This can be interpreted in terms of a unitary representation of G on $L^2(G)$: under the regular action

$$\operatorname{reg}(t)\varphi(x) = \varphi(x + t),$$

the condition of square-integrability is preserved, and in fact

$$\| \operatorname{reg}(t)\varphi \|^2 = \int_{\mathbf{R}/\mathbf{Z}} |\varphi(x + t)|^2 dx = \int_{\mathbf{R}/\mathbf{Z}} |\varphi(x)|^2 dx = \|\varphi\|^2,$$

because the Lebesgue measure is invariant under translations. Thus $L^2(G)$ is formally a unitary representation (the continuity requirement also holds; we will verify this later in a more general case). The L^2-theory of Fourier series says that the family of characters (χ_m) is an orthonormal basis of the space $L^2(G)$, and this is, quite recognizably, a suitable analogue of the decomposition of the regular representation: each of the characters χ_m appear once (a one which is the dimension of χ_m!) in $L^2(G)$. At this point it may not be entirely clear that G has no other irreducible unitary representations, but at least Schur's Lemma still applies to show that G has no *finite-dimensional* complex representation which is not one-dimensional (since G is abelian), and those are the χ_m (see Example 3.3.9). As it turns out, there is indeed no infinite-dimensional unitary representation of G, but we will see this later.

It is rather natural to explore a bit how facts about Fourier series can be interpreted in terms of the representation theory of G. Although this is quite straightforward, this brings out a few facts which are quite useful to motivate some parts of the next section, where arbitrary compact groups enter the picture.

For instance, let us consider the orthogonal projection map

$$p_m : L^2(G) \longrightarrow L^2(G)$$

onto the χ_m-isotypic component of G. Since this space is one-dimensional, with χ_m itself as a unit vector, we have simply

$$p_m(\varphi) = \langle \varphi, \chi_m \rangle \chi_m,$$

i.e., for $t\in \mathbf{R}/\mathbf{Z}$, we have

$$p_m(\varphi)(t)=\Bigl(\int_{\mathbf{R}/\mathbf{Z}}\varphi(x)e^{-2i\pi mx}dx\Bigr)e^{2i\pi mt}. \tag{5.1}$$

This seems to be a complicated way of writing the m-th Fourier coefficient of φ, which is the integral that appears here. However, we can write this as

$$\begin{aligned}p_m(\varphi)(t)&=\int_{\mathbf{R}/\mathbf{Z}}e^{2i\pi m(t-x)}\varphi(x)dx\\&=\int_{\mathbf{R}/\mathbf{Z}}e^{-2i\pi my}\varphi(y+t)dy=\int_{\mathbf{R}/\mathbf{Z}}\overline{\chi_m(y)}(\mathrm{reg}(y)\varphi)(t)dy,\end{aligned}$$

or in other words we have (formally) the formula

$$p_m=\int_{\mathbf{R}/\mathbf{Z}}\overline{\chi_m(y)}\,\mathrm{reg}(y)dy,$$

which is clearly similar to (4.20).

There is yet another instructive way to express the projection, where we consider φ as fixed and m as varying: the formula

$$p_m(\varphi)(t)=\int_{\mathbf{R}/\mathbf{Z}}\varphi(x)e^{2i\pi m(t-x)}dx$$

can be written

$$p_m(\varphi)=\varphi\star\chi_m$$

where $\cdot\star\cdot$ denotes the convolution of functions on G:

$$(\varphi_1\star\varphi_2)(t)=\int_{\mathbf{R}/\mathbf{Z}}\varphi_1(x)\varphi_2(t-x)dx$$

(when this makes sense, of course). In particular, (5.1) means that the χ_m are *eigenfunctions* of the convolution operator

$$T_\varphi:\begin{cases}L^2(G)&\longrightarrow\quad L^2(G)\\ \quad f&\mapsto\quad \varphi\star f,\end{cases}$$

(for any $\varphi\in L^2(G)$; this is a continuous linear operator on $L^2(G)$) with eigenvalues given precisely by the Fourier coefficients $\langle\varphi,\chi_m\rangle$. Note finally that this convolution operator is an intertwiner for the action of G on $L^2(G)$: this follows either from a direct computation, or from the fact that T_φ acts by scalar multiplication on the characters.

5.2. Haar measure and the regular representation of a compact group

In order to try to adapt the arguments which succeeded in the case of finite groups, and which are suggested by the example of the circle group, we see that we need first to define the analogue of the regular representation. In order for this to be a unitary representation, it seems natural to look at the space of L^2 functions on G, with respect to some "natural" measure μ. Which measure to select is dictated by the requirement that the usual action

$$\mathrm{reg}(g)\varphi(x)=\varphi(xg)$$

of the regular representation should be unitary: what is required is that

$$\int_G f(xg)d\mu(x)=\int_G f(x)d\mu(x)$$

for f integrable and $g \in G$. This property, for instance, holds for the Lebesgue measure on $\mathbf{R}/\mathbf{Z}$, or on $\mathbf{R}^d$.

It is by no means obvious that a measure μ exist with this property. Its existence is given by the following theorem:

THEOREM 5.2.1 (Existence and properties of Haar measure). *Let G be a locally compact topological group.*

(1) *There exists, up to multiplication by a scalar $\alpha \geqslant 0$, a unique Radon measure μ on G which is* right-invariant*, i.e., which is such that, for any fixed $g \in G$, we have*

$$\int_G f(xg)d\mu(x) = \int_G f(x)d\mu(x) \tag{5.2}$$

for $f \geqslant 0$ measurable and for $f \in L^1(G, d\mu)$. Such a measure, when non-zero, is called a Haar measure *on G.*

(2) *If G is compact, there exists a unique Haar measure such that $\mu(G) = 1$, which is called the* probability Haar measure *on G.*

(3) *Any Haar measure on a* compact *group G is also left-invariant, i.e., we have*

$$\int_G f(gx)d\mu(x) = \int_G f(x)d\mu(x), \tag{5.3}$$

for fixed $g \in G$, and is invariant under inversion, i.e.

$$\int_G f(x^{-1})d\mu(x) = \int_G f(x)d\mu(x),$$

both for $f \geqslant 0$ measurable or $f \in L^1(G, d\mu)$.

(4) *The support of any Haar measure on G is equal to G.*

(5) *If G is compact, let $C(G)$ be the space of continuous and bounded functions on G. Then, for any $p \geqslant 1$, the natural map $C(G) \to L^p(G, \mu)$ is injective and if $p \neq +\infty$, the space $C(G)$ is dense in $L^p(G, \mu)$ for the L^p-norm.*

REMARK 5.2.2. We have written the definition of Haar measure in terms of integral of functions. In terms of sets, μ satisfies (5.2) if, for any Borel subset of G and any $g \in G$, we have

$$\mu(B) = \mu(Bg^{-1}) = \mu(Bg)$$

(the last by applying the previous one to g^{-1}).

We will not prove this theorem in full. For many important classes of groups, it is in fact possible to write down somehow a non-zero measure μ which turns out to satisfy (5.2), and the uniqueness shows that μ is then a Haar measure.

PROOF OF (2), (3), (4). Given the existence and uniqueness statement (1), part (2) follows as soon as we check that, for a compact group G, the total measure $\mu(G)$ is finite if μ is a Haar measure. This is in fact part of the definition of a Radon measure, namely such a measure is finite on compact sets.

For (3), we fix a Haar measure μ, and then observe that for any fixed g, one can define a measure μ_g on G by

$$\int_G f(x)d\mu_g(x) = \int_G f(gx)d\mu(x),$$

and that (because left and right multiplications on G commute!) this is always right-invariant on G. Thus, by (1), there exists a non-negative real number $\Delta(g)$ such that

$$\mu_g = \Delta(g)\mu. \tag{5.4}$$

for all $y \in G$; but since G is a topological group, $Ux-1y$ is an open neighborhood of y, and therefore this relation means that y is also *not* in the support. This means that the support of μ is empty if it is not equal to all of G. Since an empty support means that the measure is identically zero, this confirms the argument at the beginning.

Finally, that continuous functions inject in L^p spaces (when G is compact) is a consequence of the fact that the support of μ is G and that bounded functions are integrable since $\mu(G) < +\infty$. The deeper fact that the continuous functions are dense in $L^p(G, \mu)$ is a general property of Radon measures. □

Remark 5.2.3. (1) It is important to remember to what extent the Haar measure is unique; if G is compact, it can be normalized uniquely by selecting as the probability Haar measure, but for a general locally compact group, there is *no* preferred Haar measure. In many applications where more than one group is involved in some problem, it becomes sometimes quite important – and sometimes delicate! – to assign suitable Haar measures to all of them...

(2) In many texts, the Haar measure is defined as a left-invariant measure (satisfying (5.3)), instead of a right-invariant one, as we did. Of course, sometimes the Haar measure as we defined it is also left-invariant (as is the case for compact groups, as we saw), and then it doesn't really matter, but otherwise one must be careful about the convention which is used (see Exercise 5.2.5 for an example of a Haar measure which is not left-invariant).

In any case, an analogue of Theorem 5.2.1 holds with left-invariant measures instead of right-invariant ones. Moreover, there is a simple link between the two: if μ is a Haar measure given by Theorem 5.2.1 (right-invariant!), defining $\Delta(g)$ by (5.4), the measure

$$d\nu(g) = \Delta(g) d\mu(g)$$

is a left-invariant measure on G (see again Exercise 5.2.5).

Example 5.2.4 (Examples of Haar measure). (1) If G is a finite group, or more generally a discrete group, then a Haar measure is given by the counting measure:

$$\mu(X) = |X|$$

for a subset $X \subset G$. For G finite, the probability Haar measure is then defined by

$$\mu(X) = \frac{|X|}{|G|}$$

for $X \subset G$, or in other words by the averaging formula

$$\int_G f(x) d\mu(x) = \frac{1}{|G|} \sum_{x \in G} f(x)$$

for $f \in C(G)$. We recognize a type of expression which was extensively used in Chapter 4!

(2) If $G = \mathbf{R}^d$, $d \geqslant 1$, or $G = (\mathbf{R}/\mathbf{Z})^d$, a Haar measure is obviously given by Lebesgue measure. (Which shows that Theorem 5.2.1 is at least as deep as the existence of the Lebesgue measure!)

(3) Let $G = \mathbf{R}^\times$. Then a Haar measure on G is given by

$$d\mu(x) = \frac{dx}{|x|}$$

in terms of the Lebesgue measure dx on $\mathbf{R}$. This can be proved by a straightforward computation using the change of variable formula for the Lebesgue measure: for $a \in \mathbf{R}^\times$, putting $y = ax$ with $dy = |a|dx$, we have

$$\int_{\mathbf{R}^\times} f(ax)\frac{dx}{|x|} = \int_{\mathbf{R}^\times} f(y)\frac{dy}{|a||y/a|} = \int_{\mathbf{R}^\times} f(y)\frac{dy}{|y|}, \tag{5.5}$$

or more conceptually using the exponential group isomorphism $\mathbf{R} \simeq \mathbf{R}^{+,\times}$ and the (obvious) fact that if

$$f \,:\, G_1 \to G_2$$

is an isomorphism of locally compact groups (i.e., a homeomorphism which is also a group isomorphism), the direct image $f_*\mu_1$ of a Haar measure on G_1 is a Haar measure on G_2.

If one restricts to $\mathbf{R}^{+,\times} =]0, +\infty[$, we get the Haar measure $x^{-1}dx$. This property explains some formulas, such as the definition

$$\Gamma(s) = \int_0^{+\infty} x^{s-1}e^{-x}dx$$

of the Gamma function: it really is an integral of $x \mapsto x^s e^{-x}$ with respect to Haar measure on $]0, +\infty[$.

(4) The example (3) can be generalized to the (genuinely non-abelian, non-compact) group $G = \mathrm{GL}_n(\mathbf{R})$ for $n \geqslant 1$: in terms of the Lebesgue measure $dx = \prod_{i,j} dx_{i,j}$ on the space $\mathrm{M}_n(\mathbf{R}) \simeq \mathbf{R}^{n^2}$ of matrices, a Haar measure on the subset G is given by

$$d\mu(x) = \frac{1}{|\det(x)|}dx. \tag{5.6}$$

This is again a simple application of the change of variable formula for multi-dimensional Lebesgue measure: the maps $x \mapsto xg$ on G extends to a diffeomorphism of $\mathrm{M}_n(\mathbf{R})$ with Jacobian $|\det(g)|$, and the result follows, formally, as in (5.5). Note that, although G is not compact (and not abelian), the left-invariance property (5.3) also holds for this Haar measure.

(5) Let $G = \mathrm{SU}_2(\mathbf{C})$. Here, and in many similar circumstances, one needs first to find a convenient parametrization to describe the Haar measure on G; typically, this means introduced finitely many continuous coordinates on G, and using a measure defined using Lebesgue measure, in terms of these coordinates.

A very convenient system of coordinates on $\mathrm{SU}_2(\mathbf{C})$ is obtained by remarking that any $g \in \mathrm{SU}_2(\mathbf{C})$ can be written uniquely

$$g = \begin{pmatrix} a & b \\ -\bar{b} & \bar{a} \end{pmatrix} \tag{5.7}$$

where a, $b \in \mathbf{C}$ are arbitrary, subject to the unique condition $|a|^2 + |b|^2 = 1$ (we leave the verification of this fact as an exercise). Using the real and imaginary parts of a and b as coordinates, we see that G is homeomorphic to the unit sphere $\mathbf{S}^3$ in $\mathbf{R}^4$.

There is an obvious measure that comes to mind on this unit sphere: the "surface" Lebesgue measure μ, which can be defined as follows in terms of the Lebesgue measure μ_4 on $\mathbf{R}^4$:

$$\mu(A) = \mu_4\Big(\Big\{x \in \mathbf{R}^4 - \{0\} \;\mid\; \|x\| \leqslant 1 \text{ and } \frac{x}{\|x\|} \in A\Big\}\Big)$$

for $A \subset \mathbf{S}^3$. This measure is natural because it is invariant under the obvious linear action of the orthogonal group $\mathrm{O}_4(\mathbf{R})$ on $\mathbf{S}^3$ (simply because so is the Lebesgue measure μ_4).

We claim that this also implies that μ is in fact a Haar measure on $\mathrm{SU}_2(\mathbf{C})$. Indeed, an element $g \in \mathrm{SU}_2(\mathbf{C})$, when acting on $\mathbf{S}^3$ by multiplication on the right, does so as the restriction of an element of $\mathrm{O}_4(\mathbf{R})$ (as an elementary computation reveals), which gives the result.

Exercise 5.2.5. We present a few additional properties and examples of Haar measures in this exercise.

(1) [Compactness and Haar measure] Show that if G is a locally compact topological group and μ is a Haar measure on G, we have $\mu(G) < +\infty$ if and only if G is compact.

Next, we will explain a slightly different proof of the left-invariance of Haar measure on a compact group, which applies to more general groups.

(2) Fix a Haar measure μ on G. Show that the function $\Delta : G \longrightarrow [0, +\infty[$ defined by (5.4) is nowhere zero, and is a continuous homomorphism $G \longrightarrow \mathbf{R}^{+,\times}$.

(3) Show that if G is compact, such a homomorphism is necessarily trivial. Show directly (without using (5.5)) that Δ is also necessarily trivial for $G = \mathrm{GL}_n(\mathbf{R})$, $n \geqslant 2$.

(4) Let G be locally compact. Show that the measure

$$d\nu(y) = \Delta(y) d\mu(y)$$

is a non-zero left-invariant measure on G.

(5) Let

$$G = \left\{ \begin{pmatrix} a & x \\ 0 & 1 \end{pmatrix} \mid a \in \mathbf{R}^{\times},\ x \in \mathbf{R}, \right\} \subset \mathrm{GL}_2(\mathbf{R}).$$

Show that, in terms of the coordinates a, b and x, the measure

$$d\mu = \frac{dadx}{|a|}$$

is a Haar measure on G. Check that it is *not* left-invariant, i.e., (5.3) does not always hold. What is the function $\Delta(g)$ in this case? What is a non-zero left-invariant measure on G?

With the Haar measure in hand, we can now define the regular representation of a locally compact group.

Proposition 5.2.6 (The regular representation). *Let G be a compact group, and let μ be a Haar measure on G. The regular action*

$$\mathrm{reg}(g)\varphi(x) = \varphi(xg)$$

is a well-defined unitary representation of G on $L^2(G, \mu)$, which is strongly continuous.

Up to isometry, this representation is independent of the choice of Haar measure, and is called "the" regular representation of G.

Similarly, the left-regular representation reg' *is defined on $L^2(G, \mu)$ by*

$$\mathrm{reg}'(g)\varphi(x) = \varphi(g^{-1}x),$$

and the two representations commute: we have $\mathrm{reg}'(g)\,\mathrm{reg}(h) = \mathrm{reg}(h)\,\mathrm{reg}'(g)$ *for all g, $h \in G$.*

Proof. Although this sounds formal – and an important goal of the theory is to set it up in such a way that it becomes formally as easy and flexible as it is in the case of finite groups –, it is important to see that there is actually some non-trivial subtleties involved.

First of all, the regular action is clearly well-defined for functions, but an element of $L^2(G, \mu)$ is an equivalence class of functions, modulo those φ which are zero μ-almost everywhere. Thus we must check that, if φ has this property, so does $\mathrm{reg}(g)\varphi$. This follows directly from the invariance of Haar measure (but it is not a triviality). Once this is settled, we check that reg is unitary using once more the invariance of μ:

$$\|\operatorname{reg}(\varphi)\|^2 = \int_G |\varphi(xg)|^2 d\mu(x) = \int_G |\varphi(x)|^2 d\mu(x) = \|\varphi\|^2.$$

What is by no means obvious is the continuity of the representation. We use the strong continuity criterion of Proposition 3.3.3 to check this.[2]

We first take $\varphi \in C(G)$, and we check the continuity at $g = 1$ of

$$g \mapsto \operatorname{reg}(g)\varphi$$

as follows: we have

$$\|\operatorname{reg}(g)\varphi - \varphi\|^2 = \int_G |\varphi(xg) - \varphi(x)|^2 d\mu(x),$$

and since $\varphi(xg) - \varphi(x) \to 0$ as $g \to 1$, for every $x \in G$, while

$$|\varphi(xg) - \varphi(x)|^2 \leqslant 4\|\varphi\|_\infty^2,$$

we see from the dominated convergence theorem that

$$\lim_{g\to 1} \|\operatorname{reg}(g)\varphi - \varphi\|^2 = 0,$$

which is the desired statement in that case.

Now if φ is arbitrary, we use the density of $C(G)$ in $L^2(G,\mu)$. Let $\varepsilon > 0$ be arbitrarily small. We find first a continuous function $\varphi_\varepsilon \in C(G)$ such that $\|\varphi - \varphi_\varepsilon\| < \varepsilon$. Then for any $g \in G$, we have

$$\begin{aligned}
\|\operatorname{reg}(g)\varphi - \varphi\| &\leqslant \|\operatorname{reg}(g)\varphi - \operatorname{reg}(g)\varphi_\varepsilon\| + \|\operatorname{reg}(g)\varphi_\varepsilon - \varphi_\varepsilon\| + \|\varphi_\varepsilon - \varphi\| \\
&= 2\|\varphi_\varepsilon - \varphi\| + \|\operatorname{reg}(g)\varphi_\varepsilon - \varphi_\varepsilon\| \\
&\leqslant 2\varepsilon + \|\operatorname{reg}(g)\varphi_\varepsilon - \varphi_\varepsilon\|.
\end{aligned}$$

By the previous case, for all g in some open neighborhood of 1 in G, we have

$$\|\operatorname{reg}(g)\varphi_\varepsilon - \varphi_\varepsilon\| < \varepsilon,$$

and hence for all such g we get

$$\|\operatorname{reg}(g)\varphi - \varphi\| \leqslant 3\varepsilon,$$

and therefore the desired strong continuity. (Note that this argument amounts to spelling out the result sketched in Exercise 3.3.12).

Finally, we observe that if we replace the Haar measure μ with $\nu = \alpha\mu$, with $\alpha > 0$, the linear map

$$\begin{cases} L^2(G,\mu) & \longrightarrow & L^2(G,\nu) \\ f & \mapsto & \alpha^{-1/2} f \end{cases}$$

is an isometry that is immediately seen to intertwine the regular representations on the two spaces. Thus the regular representation is well-defined up to isomorphism. □

Exercise 5.2.7 (Representations on L^p-spaces). Let G be a compact topological group, and μ a Haar measure on G. For any real number $p \geqslant 1$, show that there is a strongly continuous representation of G on $L^p(G,\mu)$, denoted reg_p, such that

$$\operatorname{reg}_p(g)\varphi(x) = \varphi(xg)$$

[2] Except when G is finite, it will not be continuous in the norm topology on the unitary group $\mathrm{U}(L^2(G,\mu))$.

for $\varphi \in L^p(G,\mu)$ and (almost all) $g \in G$. Check that if $\varphi \in L^2(G,\mu) \cap L^p(G,\mu)$, any L^p-translate $\mathrm{reg}_p(g)\varphi$ coincides with the corresponding translate $\mathrm{reg}(g)\varphi$ under the regular representation, in particular $\mathrm{reg}(g)\varphi \in L^2(G,\mu) \cap L^p(G,\mu)$ also. (For this reason, in the very few cases where we use the L^p-regular representation, we will omit the index p from the notation.)

Exercise 5.2.8 (The regular representation is faithful). (1) Show that the regular representation of a compact group G is faithful.

We now use this to give a simple application of representation theory to prove a general fact about compact topological groups. The goal is to show the following: for any neighborhood U of 1 in G, there exists a neighborhood $V \subset U$ which is invariant under conjugacy, i.e., such that $xVx^{-1} \subset V$ for all $x \in G$. This looks deceptively simple but it is quite tricky to prove directly (the reader may want to try!)

(2) Let $U \subset G$ be a neighborhood of 1. Show that there exists finitely many functions $f_i \in L^2(G)$ with norm 1, and $\varepsilon > 0$ such that

$$U \supset \{g \in G \mid \|\mathrm{reg}(g)f_i - f_i\| < \varepsilon \text{ for all } i\} \tag{5.8}$$

[Hint: Use (1) and the continuity of the regular representation $G \longrightarrow \mathrm{U}(L^2(\mathbf{R}))$ where the unitary group carries the strong operator topology, see Remark 3.3.4.]

(3) For a fixed index i, let $A_i \subset L^2(G)$ be the set

$$A_i = \{\mathrm{reg}(x)f_i \mid x \in G\}$$

of translates of f_i. Show that the set

$$V_i = \{g \in G \mid \|\mathrm{reg}(g)f - f\| < \varepsilon \text{ for all } f \in A_i\}$$

is conjugacy-invariant and is equal to

$$V_i = \bigcap_{x \in G} xU_ix^{-1}, \qquad U_i = \{g \in G \mid \|\mathrm{reg}(g)f_i - f_i\| < \varepsilon\}.$$

(4) To conclude, show that V_i is a neighborhood of 1, and in fact that $1 \in V_i$ and V_i is open. [Hint: Use the fact that A_i is compact.]

(5) Show that in the non-compact group $\mathrm{SL}_2(\mathbf{R})$, it is not true that any neighborhood of 1 contains a conjugacy-invariant neighborhood. Which parts of the argument above fail in that case?

What might be remembered of this exercise is the fact that the formula (5.8) shows how to use the regular representation of G to write down (a basis of) neighborhoods of 1 in G in such a way that they can be manipulated further.

Example 5.2.9 (Representations from measure-preserving actions). The regular representation of a finite group is a special case of a permutation representation. Similarly, the regular representation of a compact group can be generalized to analogues of more general permutation representations. Namely, consider a set X on which G acts (on the left), with the property that X carries a finite Radon measure ν and that G acts through measure-preserving transformations, i.e., with

$$\int_X f(g \cdot x)d\nu(x) = \int_X f(x)d\nu(x)$$

for any fixed $g \in G$ and f either $\geqslant 0$ measurable, or integrable. Then one can define a representation of G on $L^2(X,\nu)$ by

$$(g \cdot \varphi)(x) = \varphi(g^{-1} \cdot x).$$

Arguing as in the proof of Proposition 5.2.6, one sees that this is a unitary representation of G.

For a specific example, consider the group $G = \mathrm{SO}_3(\mathbf{R})$ of rotations of $\mathbf{R}^3$ acting (by matrix-vector multiplication) on the unit ball $B = \{x \in \mathbf{R}^3 \mid \|x\| \leqslant 1\}$. Taking for ν the restriction to B of Lebesgue measure on $\mathbf{R}^3$, we see that G preserves ν, and we obtain therefore a representation of $\mathrm{SO}_3(\mathbf{R})$ on $L^2(B, \nu)$.

Example 5.2.10 (Induced representations). Using the Haar measure, one can also define the proper analogue of induced representations in the setting of compact groups. Consider a compact group G, with Haar measure μ, and a compact (equivalently, closed) subgroup $K \subset G$. Given a unitary representation

$$\varrho \,:\, K \longrightarrow \mathrm{U}(H)$$

of K, one defines the induced representation $\pi = \mathrm{Ind}_K^G(\varrho)$ as follows. Define first the vector space

$$(5.9)\quad V_0 = \{f \,:\, G \longrightarrow H \mid f \text{ is continuous,} \\ \text{and } f(kg) = \varrho(k)f(g) \text{ for all } k \in K,\, g \in G\},$$

on which G acts by the regular action:

$$\pi(g)f(x) = f(xg).$$

Define an inner product on V_0 by

$$\langle f_1, f_2\rangle_0 = \int_G \langle f_1(x), f_2(x)\rangle_H d\mu(x).$$

This is well-defined (the integrand is a continuous function on G), and is positive-definite on V_0. Moreover, it is G-invariant because of the invariance of Haar measure, namely

$$\langle \pi(g)f_1, \pi(g)f_2\rangle_0 = \int_G \langle f_1(xg), f_2(xg)\rangle_H d\mu(x) = \langle f_1, f_2\rangle_0$$

for all $g \in G$.

So we almost have a unitary representation of G on V_0. But V_0 has no reason (in general) to be a Hilbert space, as completeness will usually fail. Still, one can check that π is strongly continuous on V_0 (so that it is a pre-unitary representation, as discussed in Exercise 3.3.12). Then, following the idea sketched in that exercise, we define V to be the completion of V_0 with respect to $\langle\cdot,\cdot\rangle_0$. This is a Hilbert space in which V_0 is a dense subspace, and since the $\pi(g)$ were unitary on V_0, they extend by continuity to unitary operators on V. Similarly, since the properties

$$\pi(gh) = \pi(g)\pi(h), \quad \pi(g^{-1}) = \pi(g)^{-1}$$

hold on the dense subspace $V_0 \subset V$, they are valid on all of V. The proof of the strong continuity of this representation is now obtained as in the case of the regular representation, using the fact that π is strongly continuous on V_0, unitarity, and the fact that V_0 is dense in V.

We can now see, in fact, that (just as was the case in Chapter 2, see (2.22)) the regular representation can be identified with the induced representation $\mathrm{Ind}_1^G(\mathbf{1})$. Indeed, in that case the space V_0 is the space of continuous functions on G, with its usual inner product from $L^2(G, \mu)$ and the same action as the regular representation, so that the statement amounts to the fact that $C(G)$ is dense in $L^2(G, \mu)$ for the L^2-norm.[3]

Apart from its use in describing certain representations, the most important feature of induction is Frobenius reciprocity. Looking at the proof of Proposition 2.3.6, it may seem at first to be mostly formal, and likely to extend without much ado to compact groups. However, note that part of the construction in (2.25) involves evaluating functions in the space of the induced representation at a single point, which is not well-defined in general in L^2-type spaces. Nevertheless, after proving the Peter-Weyl Theorem, we will be able to show that some important cases of Frobenius reciprocity hold (Proposition 5.4.9).

We finish this section by a discussion of unitarizability. Using integration with respect to the Haar measure, one gets the following useful fact:

Theorem 5.2.11 (Unitarizability of representations of compact groups). *Let G be a compact topological group.*

(1) *Any finite-dimensional continuous representation of G is unitarizable. As a consequence, a finite-dimensional representation of G is semisimple.*

(2) *More generally, any continuous representations*

$$\varrho : G \longrightarrow \mathrm{BGL}(H)$$

with values in the group of invertible linear maps on a Hilbert space H is unitarizable, i.e., there exists an inner product $\langle\cdot,\cdot\rangle_\varrho$ on H such that $\varrho(g) \in \mathrm{U}(H, \langle\cdot,\cdot\rangle_\varrho)$ for all $g \in G$, and such that the topology defined by this inner product is the same as the original one on H.

Proof. The first part is very easy, as in the case of finite groups: we merely use integration with respect to a Haar measure μ, instead of averaging, to construct an invariant inner product. Precisely, let

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

be a finite-dimensional complex representation of G. Fix an arbitrary inner product $\langle\cdot,\cdot\rangle_0$ on E, and define

$$\langle v, w\rangle = \int_G \langle \varrho(x)v, \varrho(x)w\rangle_0 d\mu(x).$$

The continuity of ϱ shows that the integral is well-defined (integral of a continuous bounded function); it is obviously a non-negative hermitian form product on E, and the invariance of Haar-measure shows that it is G-invariant: we have

$$\langle \varrho(g)v, \varrho(g)w\rangle = \int_G \langle \varrho(x)\varrho(g)v, \varrho(x)\varrho(g)w\rangle_0 d\mu(x) = \int_G \langle \varrho(x)v, \varrho(x)w\rangle_0 d\mu(x).$$

Moreover, note that

$$\|v\|^2 = \int_G \|\varrho(x)v\|_0^2 d\mu(x)$$

and if $v \neq 0$, this is the integral over G of a non-negative, continuous function which is non-zero since it takes the value $\|v\|_0^2 > 0$ at $x = 1$. Therefore, we have $\|v\|^2 > 0$ if $v \neq 0$, and the hermitian form is positive-definite.

Thus ϱ can be seen as a homomorphism

$$\varrho : G \longrightarrow \mathrm{U}(\langle\cdot,\cdot\rangle).$$

[3] More generally, one can give a description of the space V for an arbitrary induced representation in terms of square-integrable H-valued functions, defined in the spirit of what will be done in the next section.

As a final step, since all inner products on a finite-dimensional vector space define the same topology (and strong topology), this representation is still strongly continuous.

For (2) everything in the above goes through identically, using the given norm $\|\cdot\|$ on H instead of $\|\cdot\|_0$, except this last step: it might conceivably be the case that

$$\langle v, w\rangle_\varrho = \int_G \langle \varrho(g)v, \varrho(g)w\rangle d\mu(g)$$

defines a different topology on an infinite-dimensional Hilbert space H. However, this is not the case: for every $v \in H$, the map

$$g \mapsto \varrho(g)v$$

is continuous, and hence

$$\sup_{g\in G} \|\varrho(g)v\| < +\infty.$$

This means that the image of ϱ is "pointwise" bounded in $\mathrm{L}(H)$. The Banach-Steinhaus theorem (see, e.g., [**31**, Th. III.9]) implies that it is uniformly bounded on the unit ball, i.e., that

$$M = \sup_{g\in G} \|\varrho(g)\|_{\mathrm{L}(H)} < +\infty,$$

so that we get

$$\|\varrho(g)v\|^2 \leqslant M\|v\|^2, \qquad M^{-1}\|v\|^2 = M^{-1}\|\varrho(g^{-1})\varrho(g)v\|^2 \leqslant \|\varrho(g)v\|^2,$$

for every $g \in G$ and $v \in H$. Integrating leads to

$$M^{-1}\|v\|^2 \leqslant \langle v, v\rangle_\varrho \leqslant M\|v\|^2,$$

so that the "new" norm is topologically equivalent with the old one. □

Example 5.2.12 (Inner product on finite-dimensional representations of $\mathrm{SU}_2(\mathbf{C})$). We explain here how to compute an invariant inner product for the representation ϱ_m of $\mathrm{SU}_2(\mathbf{C})$ on the space V_m of homogeneous polynomials of degree m in $\mathbf{C}[X, Y]$.

One method is to follow the construction used in the proof of the theorem, by picking-up an inner product on V_m and averaging it over G. In order to simplify the computation, it pays to use a careful choice. One observation that may help choose wisely (here and in general) is the one in Lemma 3.3.15: unrelated unitary subrepresentations are orthogonal. This can not be applied to V_m as $\mathrm{SU}_2(\mathbf{C})$ representation (of course), but we can apply it to the restriction to a suitable subgroup. Indeed, consider the diagonal subgroup

$$K = \Bigl\{\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta}\end{pmatrix} \mid \theta \in \mathbf{R}\Bigr\},$$

so that (as in (2.37)) the space V_m decomposes as the direct sum of the subspaces $\mathbf{C}e_j$ generated by the basis vectors $e_j = X^jY^{m-j}$, $0 \leqslant j \leqslant m$, on which K acts by the one-dimensional representation

$$\varrho_m\Bigl(\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta}\end{pmatrix}\Bigr) = e^{i(2j-m)\theta}.$$

Since these are indeed distinct irreducible K-subrepresentations of V_m, it follows that for $j \neq k$, we have

$$\langle e_j, e_k\rangle = 0$$

for any invariant inner product on V_m.

Now we come back to the averaging procedure to compute the remaining inner products $\langle e_j, e_j\rangle$. However, in order to simplify again the computation, we adopt here the variant described at the end of Remark 4.1.2: it is enough to select a *non-negative* hermitian form $\langle\cdot,\cdot\rangle_0$ form on V_m, not necessarily positive-definite, provided the form defined by

$$\langle P_1, P_2\rangle = \int_{\mathrm{SU}_2(\mathbf{C})} \langle \varrho_m(g)P_1, \varrho_m(g)P_2\rangle_0 \; d\mu(g)$$

for polynomials P_1, $P_2 \in V_m$ *is* itself positive-definite. We select

$$\langle P_1, P_2\rangle_0 = P_1(1,0)\overline{P_2(1,0)},$$

and claim that this property does hold. Indeed, we get

$$\langle P, P\rangle = \int_{\mathrm{SU}_2(\mathbf{C})} |(\varrho_m(g)P)(1,0)|^2 d\mu(g) = \int_{\mathrm{SU}_2(\mathbf{C})} |P(a,b)|^2 d\mu(g),$$

so that $\|P\| = 0$ if and only if P vanishes on every $(a,b) \in \mathbf{C}^2$ which are the first row of a matrix in $\mathrm{SU}_2(\mathbf{C})$. As observed in (5.7), these are the pairs of complex numbers with $|a|^2 + |b|^2 = 1$, and by homogeneity of $P \in V_m$, it follows that in fact $P = 0$, proving that the inner product we defined is positive definite.

It remains to compute the inner products $\langle e_j, e_j\rangle$. We get

$$\langle e_j, e_j\rangle = \int_{\mathrm{SU}_2(\mathbf{C})} |a|^{2j}|b|^{2(m-j)} d\mu(g),$$

and one can quickly transform this to the beta-integral

$$\langle e_j, e_j\rangle = B(j, m-j) = \int_0^1 t^j(1-t)^{m-j}dt = \frac{1}{(m+1)\binom{m}{j}}. \tag{5.10}$$

Another argument for determining this invariant inner product is based on an a priori computation based on its known existence and its invariance property. Using the same basis as above, we find, for instance, that for all $\theta \in \mathbf{R}$ and j, k, we must have

$$\begin{aligned}\langle e_j, e_k\rangle &= \langle \begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta}\end{pmatrix} e_j, \begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta}\end{pmatrix} e_k\rangle \\ &= \langle e^{i(2j-m)\theta} e_j, e^{i(2k-m)\theta} e_k\rangle = e^{2i(j-k)\theta}\langle e_j, e_k\rangle,\end{aligned}$$

which immediately implies that e_j and e_k must be orthogonal when $j \neq k$. This means (e_j) is orthogonal basis, and only the norm $\|e_j\|^2 = \langle e_j, e_j\rangle$ need to be computed in order to conclude.

This must rely on other elements of $\mathrm{SU}_2(\mathbf{C})$ than the diagonal ones (otherwise, we would be arguing with ϱ_m restricted to the diagonal subgroup K, where any choice of $\|e_j\|^2 > 0$ gives a K-invariant inner product). Here we sketch the method in [**43**, III.2.4],[4]: consider the elements

$$g_t = \begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix}, \qquad t \in \mathbf{R},$$

of $\mathrm{SO}_2(\mathbf{R}) \subset SU_2(\mathbf{C})$, and differentiate with respect to t and evaluate at 0 the invariance formula

$$0 = \langle g_t \cdot e_j, g_t \cdot e_{j+1}\rangle.$$

One obtains the relation

$$\langle Ae_j, e_{j+1}\rangle + \langle e_j, Ae_{j+1}\rangle = 0,$$

[4] Which can be seen a simple case of studying the group $\mathrm{SU}_2(\mathbf{C})$ through its *Lie algebra*.

where A is the linear operator defined on V_m by

$$AP = \frac{d}{dt}(g_t \cdot P) \mid_{t=0} .$$

Spelling out $g_t \cdot e_j$, an elementary computation shows that

$$Ae_j = \frac{j}{2}e_{j-1} - \frac{m-j}{2}e_{j+1},$$

so that (by orthogonality of the non-diagonal inner products) we get a recurrence relation

$$(j+1)\langle e_j, e_j\rangle = (m-j)\langle e_{j+1}, e_{j+1}\rangle.$$

This determines, up to a constant scalar factor $c > 0$, the norms $\|e_j\|^2$, and indeed a quick induction leads to the formula

$$\langle e_j, e_j\rangle = cj!(m-j)!, \qquad 0 \leqslant j \leqslant m,$$

which coincides – as it should! – with (5.10) when taking $c^{-1} = (m+1)!$.

5.3. The analogue of the group algebra

It is now natural to discuss the analogue of the action of the group algebra of a finite group. However, some readers may prefer to skip to the next section, and to come back once the proof of the Peter-Weyl theorem has shown that this extension is naturally required.

The group algebra for a finite group G is the source of endomorphisms of a representation space ϱ which are linear combinations of the $\varrho(g)$. When G is compact, but possibly infinite, the group algebra (over $\mathbf{C}$) can still be defined (as finite linear combinations of symbols $[g]$, $g \in G$), but the endomorphisms it defines are not sufficient anymore. For instance, in a group like $\mathrm{SU}_2(\mathbf{C})$, the center of the group algebra is too small to create interesting intertwiners (e.g., on the regular representation), because all conjugacy classes in $\mathrm{SU}_2(\mathbf{C})$, except those of ± 1, are infinite.

It seems intuitively clear that one can solve this problem of paucity by replacing the sums defining elements of the group algebra with integrals (with respect to Haar measure). This means, that we want to consider suitable functions ψ on G and define

$$\int_G \psi(g)[g]d\mu(g),$$

in some sense. More concretely, we will consider a unitary representation ϱ, and define $\varrho(\psi)$ as the linear map

$$\varrho(\psi) = \int_G \psi(g)\varrho(g)d\mu(g). \tag{5.11}$$

This already is well-defined when ϱ is finite-dimensional: the integration can be computed coordinate-wise after selecting a basis,[5] but needs some care when ϱ is not (e.g., when G is infinite and ϱ is the regular representation), since we are then integrating a function with values in the space $\mathrm{L}(H)$ of continuous linear maps on H, for some infinite-dimensional Hilbert space H.

However, this can be defined. A natural approach would be to use the extension of Lebesgue's integration theory to Banach-space valued functions, but this is not so commonly considered in first integration courses. We use an alternate definition using "weak integrals" which is enough for our purposes and is much quicker to set-up. The

[5] One should check that the result is independent of the basis, but that is of course easy, e.g., by computing the coordinates with respect to a fixed basis.

basic outcome is: for $\psi \in L^1(G)$ (with respect to Haar measure, as usual), one can define continuous linear operators $\varrho(\psi)$, for any unitary representation ϱ of G, which behave exactly as one would expect from the formal expression (5.11) and the standard properties of integrals.

Proposition 5.3.1 (L^1-action on unitary representations). *Let G be a compact group with Haar measure μ. For any integrable function $\psi \in L^1(G)$ and any unitary representation $\varrho : G \longrightarrow \mathrm{U}(H)$, there exists a unique continuous linear operator*

$$\varrho(\psi) : H \longrightarrow H$$

such that

$$\langle \varrho(\psi)v, w\rangle = \int_G \psi(g)\langle \varrho(g)v, w\rangle d\mu(g) \tag{5.12}$$

for any vectors v, $w \in H$. This has the following properties:

(1) *For a fixed ϱ, the map*

$$\begin{cases} L^1(G) & \longrightarrow & \mathrm{L}(H) \\ \psi & \mapsto & \varrho(\psi) \end{cases}$$

is linear and continuous, with norm at most 1, i.e., $\|\varrho(\psi)v\| \leqslant \|\psi\|_{L^1}\|v\|$ for all $\psi \in L^1(G)$. If $\Phi : \varrho \longrightarrow \pi$ is a homomorphism of unitary representations, we have

$$\Phi \circ \varrho(\psi) = \pi(\psi) \circ \Phi.$$

(2) *For a fixed $\psi \in L^1(G)$, the adjoint of $\varrho(\psi)$ is given by*

$$\varrho(\psi)^* = \varrho(\check{\psi}) \tag{5.13}$$

where $\check{\psi}(g) = \overline{\psi(g^{-1})}$.

(3) *For any ψ and ϱ, and for any subrepresentation π of ϱ acting on $H_1 \subset H$, the restriction of $\varrho(\psi)$ to H_1 is given by $\pi(\psi)$. In other words, H_1 is a stable subspace of $\varrho(\psi)$.*

(4) *For any $g \in G$ and $\psi \in L^1(G)$, we have*

$$\varrho(g)\varrho(\psi) = \varrho(\mathrm{reg}'(g)\psi), \qquad \varrho(\psi)\varrho(g) = \varrho(\mathrm{reg}(g^{-1})\psi), \tag{5.14}$$

where reg *and* reg′ *represent here the right and left-regular representations acting on $L^1(G)$, as in Exercise 5.2.7. For $\psi \in L^2(G)$, this coincides with the usual regular representations.*

We will denote

$$\varrho(\psi) = \int_G \psi(x)\varrho(x)d\mu(x),$$

and for any vector $v \in H$, we also write

$$\varrho(\psi)v = \int_G \psi(x)\varrho(x)vd\mu(x),$$

which is a vector in H. The properties of the proposition can then all be seen to be formally consequences of these expressions, and can be remembered (or recovered when needed) in this way. For instance, the inequality $\|\varrho(\psi)\|_{\mathrm{L}(H)} \leqslant \|\psi\|_{L^1(G)}$ boils down to

$$\Bigl\|\int_G \psi(x)\varrho(x)d\mu(x)\Bigr\| \leqslant \int_G |\psi(x)|\|\varrho(x)\|d\mu(x) = \int_G |\psi(x)|d\mu(x)$$

which is perfectly natural (we use here that $\varrho(x)$ has norm 1 for all x).

Proof. (1) The procedure is a familiar one in Hilbert-space theory: a vector can be pinpointed (and shown to exist) by showing what its inner products with other vectors are, and these can be prescribed arbitrarily, provided only that they are linear and continuous. Precisely, given $\psi \in L^1(G)$, a unitary representation ϱ on H and a vector $v \in H$, define

$$\ell_{\psi,v} \begin{cases} H \longrightarrow \mathbf{C} \\ w \mapsto \displaystyle\int_G \psi(g)\langle w, \varrho(g)v\rangle d\mu(g). \end{cases}$$

This is well-defined because ψ is integrable and the factor

$$|\langle w, \varrho(g)v\rangle| \leqslant \|w\|\|v\|$$

is bounded. It is also continuous, by strong continuity of ϱ, hence measurable. Moreover, the map $\ell_{\psi,v}$ is clearly a linear form on H, and it is continuous, since the same inequality leads to

$$|\ell_{\psi,v}(w)| \leqslant C\|w\|, \qquad C = \|v\|\|\psi\|_{L^1(G)}$$

for all $w \in H$. According to the Riesz representation theorem for Hilbert spaces, there exists therefore a unique vector, which we denote $\varrho(\psi)v$, such that

$$\ell_{\psi,v}(w) = \langle w, \varrho(\psi)v\rangle$$

for all $w \in H$. Taking the conjugate, we obtain (5.12), and the uniqueness is then also achieved.

From this, the remaining properties are quite easily checked. For (1), the linearity (both the linearity of $\varrho(\psi)$ as a map on H, and then the linearity as a function of ψ) is deduced in a very standard way from the uniqueness and the linearity of $\ell_{\psi,v}$ as function of v and ψ.

Riesz's Theorem also implies that $\|\varrho(\psi)v\| = \|\ell_{\psi,v}\|$, which is bounded by the constant $C = \|v\|\|\psi\|_{L^1}$ above. This inequality

$$\|\varrho(\psi)v\| \leqslant \|\psi\|_{L^1}\|v\|$$

shows that $\varrho(\psi)$ is continuous for a fixed ψ, and also that it is continuous as a map $L^1(G) \longrightarrow \mathrm{L}(H)$, with norm at most 1 as claimed.

(2) We leave this as an exercise.

(3) plays an important role later on, so we give the argument: let $v \in H_1$ be a vector in the stable subspace; we first check that $\varrho(\psi)v$ is also in H_1, using orthogonality. If $w \in H_1^\perp$ is orthogonal to H_1, we have

$$\langle \varrho(\psi)v, w\rangle = \int_G \psi(g)\langle \varrho(g)v, w\rangle d\mu(g) = 0,$$

i.e., $\varrho(\psi)v \in (H_1^\perp)^\perp = H_1$, as desired. But then, $\varrho(\psi)v$ is determined by its inner products with vectors $w \in H_1$, and then we have

$$\langle \pi(\psi)v, w\rangle = \int_G \psi(g)\langle \pi(g)v, w\rangle d\mu(g)$$

by definition, which – since π "is" simply ϱ restricted to H_1 – is equal to

$$\int_G \psi(g)\langle \varrho(g)v, w\rangle d\mu(g) = \langle \varrho(\psi)v, w\rangle$$

as expected.

(4) These formulas replace formal computations based on the invariance of the Haar measure under translation. We only write down one of these: for $g \in G$ and $\psi \in L^1(G)$, we have

$$\varrho(g)\varrho(\psi) = \varrho(g)\int_G \psi(x)\varrho(x)d\mu(x) = \int_G \psi(x)\varrho(gx)d\mu(x) = \int_G \psi(g^{-1}y)\varrho(y)d\mu(y)$$

which "is" is the first formula. We leave the other to the reader, as we leave her the task of translating it into a formal argument □

Exercise 5.3.2 (More general integrals). Many variants of this construction are possible. For instance, let H be a separable Hilbert space (so there is a countable subset of H which is dense in H). For any function

$$f : G \longrightarrow H$$

which is "weakly measurable", in the sense that for every $w \in H$, the function

$$g \mapsto \langle f(g), w\rangle$$

is measurable on G, and which has bounded values (in the norm of H), show how to define the integrals

$$\int_G f(g)d\mu(g) \in H$$

and show that this construction is linear with respect to f, and satisfies

$$\Bigl\|\int_G f(g)d\mu(g)\Bigr\| \leqslant \int_G \|f(g)\|d\mu(g) \tag{5.15}$$

(you will first have to show that $g \mapsto \|f(g)\|$ is integrable). Moreover, for a unitary representation ϱ on H and for any $\psi \in L^1(G)$ and $v \in H$, show that

$$\varrho(\psi)v = \int_G f(g)d\mu(g)$$

for $f(g) = \psi(g)\varrho(g)v$.

Exercise 5.3.3 (Intertwiners from L^1-action). Let G be a compact group and μ a Haar measure on G.

(1) For an integrable function φ on G, explain how to define the fact that φ is a class function (i.e., formally, $\varphi(xyx^{-1}) = \varphi(y)$ for all x and y in G).

(2) Let φ be an integrable class function on G. For any unitary representation ϱ of G, show that the operator

$$\int_G \varphi(x)\varrho(x)d\mu(x)$$

is in $\mathrm{Hom}_G(\varrho, \varrho)$, i.e., is an intertwiner.

Exercise 5.3.4 (Convolution). Let G be a compact group and μ a Haar measure on G. For any functions φ, $\psi \in L^2(G)$, show that

$$\mathrm{reg}'(\varphi)\psi = \varphi \star \psi,$$

where the convolution $\varphi \star \psi$ is defined by

$$(\varphi \star \psi)(g) = \int_G \varphi(x)\psi(x^{-1}g)d\mu(x).$$

Prove also that

$$(\varphi \star \psi)(g) = \langle \varphi, \mathrm{reg}'(g)\check{\psi}\rangle,$$

where $\check{\psi}$ is given by (5.13), and deduce that $\varphi \star \psi$ is continuous.

The formulas (5.14) express the link between the action of L^1-functions on H (given by the "implicit" definition (5.12) using inner products) and the original representation ϱ. The following is another important relation; in fact, it shows how to recover the original representation operators $\varrho(g)$ starting from the collection of linear maps $\varrho(\psi)$. This is trickier than for finite groups, because there is (in the infinite case) no integrable function supported only at a single point $g \in G$.

Proposition 5.3.5 (L^1-approximation of representation operators). *Let G be a compact topological group, and let $\varrho : G \longrightarrow \mathrm{U}(H)$ be a unitary representation of G. Fix $g \in G$, and for any open neighborhood U of g in G, write ψ_U for the characteristic function of U, normalized so that $\|\psi_U\| = 1$, i.e., define*

$$\psi_U(x) = \begin{cases} \dfrac{1}{\mu(U)} & \text{if } x \in U \\ 0 & \text{otherwise.} \end{cases}$$

Then we have

$$\lim_{U \downarrow g} \varrho(\psi_U) = \varrho(g),$$

where the limit is a limit "as U tends to g", taken in the strong topology in $\mathrm{L}(H)$, i.e., it should be interpreted as follows: for any $v \in H$, and for any $\varepsilon > 0$, there exists an open neighborhood V of g such that

$$\|\varrho(\psi_U)v - \varrho(g)v\| < \varepsilon, \tag{5.16}$$

for any $U \subset V$.

Proof. This is a simple analogue of classical "approximation by convolution" results in integration theory. We use (5.12), and the fact that the integral of ψ_U is one, to write

$$\langle \varrho(\psi_U)v - \varrho(g)v, w\rangle = \int_G \psi_U(x)\langle \varrho(x)v - \varrho(g)v, w\rangle d\mu(x)$$

for any U and any $w \in H$. Therefore

$$\begin{aligned}|\langle \varrho(\psi_U)v - \varrho(g)v, w\rangle| &\leqslant \|w\| \int_G \psi_U(x)\|\varrho(x)v - \varrho(g)v\| d\mu(x) \\ &\leqslant \|w\| \sup_{x\in U} \|\varrho(x)v - \varrho(g)v\|\end{aligned}$$

for all $w \in H$. This implies that

$$\|\varrho(\psi_U)v - \varrho(g)v\| \leqslant \sup_{x\in U} \|\varrho(x)v - \varrho(g)v\|$$

(alternatively, this should be thought as an application of the inequality (5.15) for the integral

$$\varrho(\psi_U)v - \varrho(g)v = \int_G \psi_U(x)(\varrho(x)v - \varrho(g)v)d\mu(x)$$

as defined in Exercise 5.3.2.)

But now, the strong continuity of the representation ϱ means that $x \mapsto \|\varrho(x)v - \varrho(g)v\|$ is a continuous function on G taking the value 0 at $x = g$. Hence, for any $\varepsilon > 0$, we can find some neighborhood V of g such that

$$\|\varrho(x)v - \varrho(g)v\| < \varepsilon$$

for all $x \in V$. This choice of V gives the desired inequality (5.16). □

5.4. The Peter-Weyl theorem

Once we have the regular representation reg of a compact group G, we attack the study of unitary representations of the group by attempting to decompose it into irreducibles. This is done by the Peter-Weyl theorem, from which all fundamental facts about the representations of general compact groups follow.

THEOREM 5.4.1 (Peter-Weyl). *Let G be a compact topological group with probability Haar measure μ. Then the regular representation of G on $L^2(G, \mu)$ decomposes as a Hilbert space direct sum*[6]

$$L^2(G, \mu) = \bigoplus_{\varrho} M(\varrho) \tag{5.17}$$

of isotypic components of the finite-dimensional *irreducible unitary representations of G, each $M(\varrho)$ being isomorphic to a direct sum of* $\dim(\varrho)$ *copies of ϱ.*

Although this result addresses only the properties of the regular representation, it should not be surprising, in view of the importance of the latter in the case of finite groups, that it leads to results for arbitrary unitary representations:

COROLLARY 5.4.2 (Decomposition of unitary representations). *Let G be a compact topological group with probability Haar measure μ.*

(1) *Any irreducible unitary representation of G is finite-dimensional.*

(2) *Any unitary representation of G is isomorphic to a Hilbert direct sum of irreducible subrepresentations.*

We start with the proof of the Peter-Weyl theorem. As usual, we attempt to motivate the arguments, instead of trying to present the shortest proof possible.

The first observation we can make is that, essentially, we already know how to obtain the inclusion

$$\bigoplus_{\varrho} M(\varrho) \subset L^2(G, \mu) \tag{5.18}$$

(where the direct sum is orthogonal) as well as the fact that $M(\varrho)$ is, for any finite-dimensional irreducible unitary representation, isomorphic to a direct sum of $\dim(\varrho)$ copies of ϱ.

Indeed, we can follow the method of Section 2.7.3, and in particular Theorem 2.7.24, to embed irreducible finite-dimensional representations in $L^2(G, \mu)$. Given an irreducible unitary representation

$$\varrho \,:\, G \longrightarrow \mathrm{U}(H)$$

of G, we see that the unitary matrix coefficients

$$f_{v,w} \,:\, g \mapsto \langle \varrho(g)v, w\rangle$$

are bounded functions on G, by the Cauchy-Schwarz inequality, and are continuous, so that $f_{v,w} \in L^2(G)$ for all v, $w \in H$ (this is (5) in Theorem 5.2.1, which also shows that distinct matrix coefficients are distinct in $L^2(G)$). A formal argument (as in Theorem 2.7.24) implies that, for a fixed $w \in H$, the map

$$v \mapsto f_{v,w}$$

[6] Recall that we defined a orthogonal direct sum of unitary representations in Section 3.3.

is an intertwiner of ϱ and reg. If $w \neq 0$, Schur's Lemma implies that this map is injective, because it is then non-zero: we have $f_{w,w}(1) = \|w\|^2 \neq 0$, and the continuity of $f_{w,w}$ shows that it is a non-zero element of $L^2(G)$.

Now we assume in addition that ϱ is finite-dimensional. In that case, the image of $v \mapsto f_{v,w}$ is a closed subspace of $L^2(G)$ (since any finite-dimensional subspace of a Banach space is closed) and hence it is a subrepresentation of reg which is isomorphic to ϱ. Varying w, again as in Theorem 2.7.24, we see furthermore that reg contains a subrepresentation isomorphic to a direct sum of $\dim(\varrho)$ copies of ϱ.

If we let ϱ vary among non-isomorphic finite-dimensional unitary representations, we also see that, by Lemma 3.3.15, the corresponding spaces of matrix coefficients are orthogonal.

So to finish the proof of the first inclusion (5.18), we should only check that the ϱ-isotypic component coincides with the space of matrix coefficients of ϱ. We can expect this to hold, of course, from the case of finite groups. Indeed, this is true, but this time the argument in the proof of Theorem 2.7.24 needs some care before it can be applied, as it relies on the linear form $\delta : \varphi \mapsto \varphi(1)$ which is not necessarily continuous on a subspace of $L^2(G)$ (and it is certainly not continuous on all of $L^2(G)$, if G is infinite).

Still, δ is well-defined on any space consisting of continuous functions. The following lemma will then prove to be ad-hoc:

Lemma 5.4.3. *Let G be a compact group and let $E \subset L^2(G)$ be a* finite-dimensional *subrepresentation of the regular representation. Then E is (the image of) a space of continuous functions.*

Assuming this, let $E \subset L^2(G)$ be any subrepresentation isomorphic to ϱ. We can view E, by the lemma, as a subspace of $C(G)$. Then the linear form

$$\delta : \varphi \mapsto \varphi(1)$$

is well-defined on E. Using the inner product induced on E by the L^2-inner product, it follows that there exists a unique $\psi_0 \in E$ such that

$$\delta(\varphi) = \langle \varphi, \psi_0 \rangle$$

for all $\varphi \in E$. We conclude as in the proof of Theorem 2.7.24: for all $\varphi \in E$ and $x \in G$, we have

$$\varphi(x) = \delta(\mathrm{reg}(x)\varphi) = \langle \mathrm{reg}(x)\varphi, \psi_0 \rangle = f_{\varphi,\psi_0}(x),$$

so that $\varphi = f_{\varphi,\psi_0}$. This shows that E is contained in the space of matrix coefficients of ϱ.

Before going to the converse inclusion, we must prove Lemma 5.4.3:

Proof of Lemma 5.4.3. The basic idea is that continuous functions can be obtained by averaging integrable functions, and that averaging translates of a given $\varphi \in E$ (under the regular representation) leads to another function in E. This way we will show that $E \cap C(G)$ is dense in E (for the L^2-norm), and since $\dim E < +\infty$, this implies that $E \cap C(G) = E$, which is what we want.

Thus, given $\varphi \in E$, consider functions of the type

$$\vartheta(g) = \int_G \psi(x)\varphi(gx)d\mu(x)$$

where $\psi \in L^2(G)$. If we write this as

$$\vartheta(g) = \langle \mathrm{reg}'(g^{-1})\varphi, \overline{\psi} \rangle,$$

the strong continuity of the left-regular representation shows that ϑ is continuous. But we can also write this expression, as

$$\vartheta = \int_G \psi(x)\,\mathrm{reg}(x)\varphi \; d\mu(x) = \mathrm{reg}(\psi)\varphi,$$

using the action of L^1 functions defined in Proposition 5.3.1 (since G is compact, any square-integrable function is also integrable). This shows (using part (3) of the proposition) that $\vartheta \in E \cap C(G)$. Finally, by Proposition 5.3.5 applied to $g = 1$, for any $\varepsilon > 0$, we can find $\psi \in L^1(G)$ such that

$$\|\varphi - \vartheta\| = \|\,\mathrm{reg}(1)\varphi - \mathrm{reg}(\psi)\varphi\| \leqslant \varepsilon,$$

and we are done. □

We have now proved (5.18), and must consider the converse. The problem is that, for all we know, the set of finite-dimensional irreducible unitary representations of G might be reduced to the trivial one! In other words, to prove the reverse inclusion, we need a way to construct or show the existence of finite-dimensional representations of G.

The following exercise shows an "easy" way, which applies to certain important groups (like $\mathrm{U}_n(\mathbf{C})$ for instance).

Exercise 5.4.4 (Groups with a finite-dimensional faithful representation). Let G be a compact topological group which is a closed subgroup of $\mathrm{GL}_n(\mathbf{C})$ for some $n \geqslant 1$.

Show that the linear span of matrix coefficients of finite-dimensional irreducible representations of G is a dense subspace of $C(G)$, using the Stone-Weierstrass theorem (we recall the statement of the latter in Section A.3). Deduce the Peter-Weyl Theorem from this.

The class of groups covered by this exercise is restricted (in the next chapter, we describe another characterization of these groups and give examples of compact groups which are not of this type). We now deal with the general case, by proving that the direct sum

$$\bigoplus_{\varrho} M(\varrho)$$

is dense in $L^2(G)$. Equivalently, using Hilbert space theory, we must show that if $\varphi \in L^2(G)$ is non-zero, it is *not* orthogonal to all $M(\varrho)$.

The basic motivation for what follows comes from looking at the case of the circle group: when ϱ varies, we expect the projection onto $M(\varrho)$ to be given by a convolution operator that commutes with the regular representation; if we can find a non-zero eigenvalue with finite multiplicity of this convolution operator, this will correspond to a non-trivial projection.

The details are a bit more involved because the group G is not necessarily abelian. We exploit the fact that $M(\varrho)$ is stable under the left-regular representation reg': if $\varphi \perp M(\varrho)$, we have

$$\langle \varphi, \mathrm{reg}'(g) f\rangle = 0$$

for all $g \in G$ and $f \in M(\varrho)$. As a function of g, this is the convolution $\varphi \star \check{f}$ (see Exercise 5.3.4). If we further note that, for a basic matrix coefficient $f(g) = \langle \varrho(g)v, w\rangle$, we have

$$\check{f}(g) = \overline{\langle \varrho(g^{-1})v, w\rangle} = \langle \varrho(g)w, v\rangle$$

which is also a matrix coefficient, we can deduce that $\varphi \perp M(\varrho)$ implies that $\varphi \star f = 0$ for all $f \in M(\varrho)$.

It is therefore sufficient to prove that, for some finite-dimensional subrepresentation $E \subset L^2(G)$, the convolution operator

$$T_\varphi : f \mapsto \varphi \star f$$

is non-zero on E. For an arbitrary operator, this might be very tricky, but by Exercise 5.3.4 again, we have also

$$T_\varphi(f) = \mathrm{reg}'(\varphi)f,$$

so that T_φ is an intertwiner of the regular representation (Exercise 5.3.3). As such, by Schur's Lemma, it acts by a scalar on any finite-dimensional subrepresentation. The question is whether any of these scalars is non-zero, i.e., whether T_φ has a non-zero eigenvalue when acting on these finite-dimensional subspaces. Now here comes the trick: we basically want to claim that the convolution form of T_φ, *abstractly*, implies that it has a non-zero eigenspace $\ker(T_\varphi - \lambda)$, with non-zero eigenvalue λ, of finite dimension. If that is the case, then $\ker(T_\varphi - \lambda)$ is a finite-dimensional subrepresentation on which T_φ is a non-zero scalar, and we deduce that φ is not orthogonal to it.

To actually implement this we must change the operator to obtain a better-behaved one, more precisely a self-adjoint operator. We form the function $\psi = \check{\varphi} \star \varphi$, and consider the convolution operator

$$T_\psi = \mathrm{reg}'(\psi) = \mathrm{reg}'(\varphi)^* \,\mathrm{reg}'(\varphi).$$

Since the convolution product is associative, we see that $\ker(T_\varphi) \subset \ker(T_\psi)$, and hence the previous reasoning applies to T_ψ also: T_ψ intertwines the regular representation with itself, and if there exists a finite-dimensional subrepresentation E such that T_ψ acts as a non-zero scalar on E, the function φ is not orthogonal to E.

But now T_ψ is self-adjoint, and even non-negative since

$$\langle T_\psi f, f\rangle = \|T_\varphi f\|^2 \geqslant 0$$

for $f \in L^2(G)$. Moreover, writing

$$T_\psi f(x) = \mathrm{reg}'(\psi)f(x) = \int_G \psi(y) f(y^{-1}x) d\mu(y) = \int_G \psi(xy^{-1}) f(y) d\mu(y)$$

(by invariance of Haar measure), we see that T_ψ is an integral Hilbert-Schmidt operator on G with kernel

$$k(x, y) = \psi(xy^{-1})$$

(as defined in Proposition A.2.4); this kernel is in $L^2(G \times G, \mu \times \mu)$ (again by invariant of Haar measure, its L^2-norm is $\|\psi\|$). By Proposition A.2.4, it is therefore a compact operator.

The operator T_ψ is also non-zero, because ψ is continuous (Exercise 5.3.4 again) and $\psi(1) = \|\varphi\|^2 \neq 0$ (by assumption), and because by taking f to be a normalized characteristic function of a small enough neighborhood of 1, we have $\psi \star f$ close to ψ, hence non-zero (this is Proposition 5.3.5, applied to the left-regular representation). So we can apply the spectral theorem (Theorem A.2.3, or the minimal version stated in Proposition A.2.5) to deduce that T_ψ has a non-zero eigenvalue λ with finite-dimensional eigenspace $\ker(T_\psi - \lambda)$, as desired.

Remark 5.4.5. At the end of Section A.2, we prove – essentially from scratch – the minimal part of the spectral theorem which suffices for this argument, in the case when the space $L^2(G, \mu)$ is separable (which is true, for instance, whenever the topology of G is defined by a metric and thus for most groups of practical interest; see [**31**, Pb. IV.43] for this.)

Exercise 5.4.6 (Another argument). Let G be a compact topological group with probability Haar measure μ.

(1) Show directly that, for any $g \neq 1$, there exists a finite-dimensional unitary representation ϱ of G such that $\varrho(g) \neq 1$. (We also state this fact formally in Corollary 5.4.8.)[Hint: One can also use compact operators for this purpose.]

(2) Use this and the Stone-Weierstrass Theorem (Theorem A.3.1) to give a proof of the Peter-Weyl theorem in the general case. (See Exercise 5.4.4.)

We now come to the proof of Corollary 5.4.2. This also requires the construction of some finite-dimensional subrepresentations. The following lemma is therefore clearly useful:

Lemma 5.4.7. *Let G be a compact topological group, and let*

$$\varrho \,:\, G \longrightarrow \mathrm{U}(H)$$

be any non-zero unitary representation of G. Then ϱ contains an irreducible finite-dimensional subrepresentation.

Proof. It suffices to find a non-zero finite-dimensional subrepresentation of H, since it will in turn contain an irreducible one. But we can not hope to construct the desired subrepresentation using kernels or eigenspaces of intertwiners this time (think of ϱ being an infinite orthogonal direct sum of copies of the same irreducible representation). Instead we bootstrap the knowledge we just acquired of the regular representation, and use the following remark: if $E \subset L^2(G)$ is a finite-dimensional subrepresentation of the left-regular representation and $v \in H$, then the image

$$F = \{\mathrm{reg}'(\varphi)v \mid \varphi \in E\}$$

of the action of E on v is a subrepresentation of H. Indeed, by (5.14), we have

$$\varrho(g)\varrho(\varphi)v = \varrho(\mathrm{reg}'(g)\varphi)v \in F$$

for all $\varphi \in E$ and $g \in G$. Obviously, the subspace F is a quotient of E (by the obvious surjection $\varphi \mapsto \varrho(\varphi)v$), and hence we will be done if we can ensure that $F \neq 0$.

To do this, we fix any $v \neq 0$, and basically use the fact that $\varrho(1)v = v \neq 0$ is "almost" in F. So we approximate $\varrho(1)$ using the L^2-action: first of all, by Proposition 5.3.5, we can find $\psi \in L^2(G)$ such that $\varrho(\psi)v$ is arbitrarily close to $\varrho(1)v$, in particular, we can ensure that $\varrho(\psi)v \neq 0$. Further, using the Peter-Weyl Theorem, we can find a $\psi_1 \in L^2(G)$ which is a (finite) linear combination of matrix coefficients of finite-dimensional representations and which approximates ψ arbitrarily closely. Since (Proposition 5.3.1, (1)) we have

$$\|\varrho(\psi) - \varrho(\psi_1)\|_{\mathrm{L}(H)} \leqslant \|\psi - \psi_1\|_{L^1} \leqslant \|\psi - \psi_1\|_{L^2},$$

we get

$$\|\varrho(\psi_1)v\|_H \geqslant \|\varrho(\psi)v\|_H - \|\psi - \psi_1\|_{L^2}\|v\|_H,$$

which will be > 0 if the approximation ψ_1 is suitably chosen. Now take for E the finite direct sum of the spaces $M(\varrho)$ for the ϱ which occur in the expression of ψ_1 as a combination of matrix coefficients: we have $\psi_1 \in E$, and E is a subrepresentation of the left-regular representation with $F \neq 0$ since it contains $\varrho(\psi_1)v \neq 0$. □

Proof of Corollary 5.4.2. First of all, Lemma 5.4.7 shows by contraposition that all irreducible representations must be finite-dimensional (if ϱ is infinite-dimensional, the statement shows it is not irreducible). Then we also see that the "complete reducibility" must be true: if reducibility failed, the (non-zero) orthogonal complement of a "maximal" completely-reducible subrepresentation could not satisfy the conclusion of the lemma. To be rigorous, this reasoning can be expressed using Zorn's Lemma. We give the details, though the reader may certainly think that this is quite obvious (or may rather write his own proof). Let $\varrho : G \longrightarrow \mathrm{U}(H)$ be a unitary representation; consider the set

$$\mathcal{O} = \{(I, (H_i)_{i\in I})\}$$

where I is an arbitrary index set, and $H_i \subset H$ are pairwise orthogonal finite-dimensional irreducible subrepresentations of G. This set is not empty, by Lemma 5.4.7. We order it by inclusion: $(I, (H_i)) \leqslant (J, (H'_j))$ if and only if $I \subset J$ and $H'_i = H_i$ for $i \in I \subset J$. This complicated-looking ordered set is set up so that it is very easy to see that every totally ordered subset $\mathcal{P}$ has an upper bound.[7] By Zorn's Lemma, we can find a maximal element $(I, (H_i))$ in $\mathcal{O}$. Then we claim that the subspace

$$H' = \bigoplus_{i\in I}^{\perp} H_i$$

is in fact equal to H, which is then exhibited as a Hilbert orthogonal sum of finite-dimensional subspaces. Indeed, consider $H'' = (H')^{\perp} \subset H$. If this is non-zero, H'' contains a finite-dimensional subrepresentation, say H_0, again by Lemma 5.4.7. But then (assuming the index 0 is not in I...) we have

$$(I \cup \{0\}, (H_i, H_0)) \in \mathcal{O},$$

contradicting the maximality of $(I, (H_i))$. Thus $H'' = 0$, which means that $H' = H$ as claimed. □

We finally deduce some further corollaries of the Peter-Weyl theorem:

Corollary 5.4.8 (Separating points and completeness criteria). *Let G be a compact topological group with probability Haar measure μ.*

(1) *If $g \neq 1$ is a non-trivial element of G, there exists a finite-dimensional unitary, indeed irreducible, representation ϱ of G such that $\varrho(g) \neq 1$.*

(2) *Suppose $\mathcal{C}$ is a given set of finite-dimensional irreducible unitary representations of G. Then $\mathcal{C}$ is complete, i.e., contains all irreducible representations of G, if and only if the linear span M of the matrix coefficients of representations $\varrho \in \mathcal{C}$ is dense in $L^2(G, \mu)$, or equivalently if M is dense in $C(G)$ for the L^∞-norm.*

(3) *Suppose $\mathcal{C}$ is a given set of finite-dimensional irreducible unitary representations of G; then $\mathcal{C}$ is complete if and only if we have the Parseval formula*

$$\|\varphi\|^2 = \sum_{\varrho\in\mathcal{C}} \|p_\varrho(\varphi)\|^2$$

for all $\varphi \in L^2(G)$, where p_ϱ is the orthogonal projection on the isotypic component $M(\varrho)$.

As a matter of fact, the statement of the Peter-Weyl Theorem in the original paper is that (3) holds (and the statement in Pontryagin's version a few years later was the density of M in $C(G)$!

Proof. (1) The regular representation is faithful (Exercise 5.2.8), so for $g \neq 1$, the operator $\mathrm{reg}(g)$ is not the identity. However, by the Peter-Weyl Theorem, $\mathrm{reg}(g)$ is the direct sum of the operators obtained by restriction to each $M(\varrho)$, which are just direct

[7] For the more natural set $\mathcal{O}'$ of all "completely reducible subrepresentations" of H, ordered by inclusion, checking this is more painful, because one is not keeping track of consistent decompositions of the subspaces to use to construct an upper bound.

sums of $\dim(\varrho)$ copies of $\varrho(g)$. Hence some at least among the $\varrho(g)$ must be distinct from the identity.

The statement of (2) concerning the L^2-norm is an immediate consequence of the Peter-Weyl Theorem and Hilbert space theory: if M is not dense in $L^2(G,\mu)$, the orthogonal complement of its closure is a non-zero subrepresentation of the regular representation, which must therefore contain some irreducible subrepresentation π. Because the definition of M means that it is spanned by the $M(\varrho)$ where ϱ ranges over $\mathcal{C}$, we must have $\pi \notin \mathcal{C}$.

Similarly, (3) is a direct translation of the Peter-Weyl Theorem using the theory of Hilbert spaces.

For the part of (2) involving continuous functions, we recall that M is indeed a subspace of $C(G)$. We must show that M is dense in $C(G)$ for the L^∞-norm. This can be thought of as an analogue of the classical Weierstrass approximation theorem for trigonometric polynomials (which, indeed, corresponds to $G = \mathbf{S}^1$), and one can indeed use (1) and the Stone-Weierstrass Theorem (this is the content of Exercise 5.4.6). □

Our last result in this section is a special case of Frobenius reciprocity for induced representations.

Proposition 5.4.9 (Frobenius reciprocity for compact groups). *Let G be a compact topological group with probability Haar measure μ, and let $K \subset G$ be a compact subgroup. For any* finite-dimensional *unitary representations*

$$\varrho_1 \,:\, G \longrightarrow \mathrm{U}(H_1), \qquad \varrho_2 \,:\, K \longrightarrow \mathrm{U}(H_2),$$

we have a natural isomorphism

$$\mathrm{Hom}_G(\varrho_1, \mathrm{Ind}_K^G(\varrho_2)) \simeq \mathrm{Hom}_K(\mathrm{Res}_K^G(\varrho_1), \varrho_2).$$

Proof. We leave it to the reader to check that the proof of Proposition 2.3.6 can carry through formally unchanged, provided one proves first that the image of any intertwiner

$$\Phi \,:\, \varrho_1 \longrightarrow \mathrm{Ind}_K^G(\varrho_2)$$

lies in the (image of) the subspace V_0 of continuous functions used in the definition of induced representations (see (5.9); note that all intertwiners constructed from right to left by (2.26) have this property, because of the strong continuity of unitary representations, so that Frobenius reciprocity can only hold when this property is true.)

But since ϱ_1 is finite-dimensional, so is the image of Φ, and thus the statement is an analogue of Lemma 5.4.3. To reduce to that case, note the isomorphism

$$\mathrm{Ind}_K^G(\varrho_2) \simeq L^2(G) \oplus \cdots \oplus L^2(G)$$

of K-representations (with $\dim(\varrho_2)$ copies on the right) given by mapping a function $f \in V$ to its components with respect to a fixed orthonormal basis (e_i) of H_2, i.e., we have

$$f(x) = \sum_i f_i(x) e_i$$

for $f \in V$. By Lemma 5.4.3, each f_i is continuous for $f \in \mathrm{Im}(\Phi)$, and hence $\mathrm{Im}(\Phi) \subset V_0$. □

Exercise 5.4.10. Which of the other properties of induction can you establish for compact groups?

5.5. Characters and matrix coefficients for compact groups

The reward for carefully selecting the conditions defining unitary representations of compact groups and proving the analytic side of the Peter-Weyl Theorem is that, with this in hand, character theory becomes available, and is just as remarkably efficient as in the case of finite groups (but of course it is also restricted to finite-dimensional representations).

We summarize the most important properties, using as before the notation $\hat{G}$ for the set of irreducible unitary representations of G.

Theorem 5.5.1 (Character theory). *Let G be a compact topological group, and let μ be the probability Haar measure on G.*

(1) *The characters of irreducible unitary representations of G are continuous functions on G, which form an orthonormal basis of the space $L^2(G^\sharp)$ of conjugacy-invariant functions on G. In particular, a set $\mathcal{C}$ of finite-dimensional irreducible representations of G is complete, in the sense of Corollary 5.4.8, if and only if the linear combinations of characters of representations in $\mathcal{C}$ are dense in $L^2(G^\sharp)$.*

(2) *For any irreducible unitary representation $\pi \in \hat{G}$ of G, and any unitary representation*

$$\varrho : G \longrightarrow \mathrm{U}(H)$$

of G, the orthogonal projection map onto the π-isotypic component of H is given by

$$\Phi_\pi = (\dim \pi) \int_G \overline{\chi_\pi(g)} \varrho(g) d\mu(g),$$

using the L^1-action of Proposition 5.3.1.

In particular, if ϱ is finite-dimensional, the dimension of the space ϱ^G of G-invariant vectors in H is given by the average over G of the values of the character of ϱ, i.e.,

$$\dim \varrho^G = \int_G \chi_\varrho(g) d\mu(g).$$

(3) *A unitary finite-dimensional representation ϱ of G is irreducible if and only if*

$$\int_G |\chi_\varrho(g)|^2 d\mu(g) = 1.$$

More generally, if ϱ_1 and ϱ_2 are finite-dimensional *unitary representation of G, we have*

$$\langle \chi_{\varrho_1}, \chi_{\varrho_2} \rangle = \sum_{\pi \in \hat{G}} n_\pi(\varrho_1) n_\pi(\varrho_2), \tag{5.19}$$

where $n_\pi(\varrho) = \langle \chi_\varrho, \chi_\pi \rangle$ is the multiplicity of π in ϱ.

The reader should definitely try her hand at checking all these facts, without looking at the proof, since they are analogues of things we know for finite groups. For the sake of variety, we use slightly different arguments (which can also be applied to finite groups.) First we compute the inner products of matrix coefficients:

Lemma 5.5.2. *Let G be a compact group with probability Haar measure μ.*

(1) *If π_1 and π_2 are non-isomorphic irreducible unitary representations of G, any two matrix coefficients of π_1 and π_2 are orthogonal in $L^2(G)$.*

(2) *If* $\pi : G \longrightarrow \mathrm{U}(H)$ *is an irreducible unitary representation of* G *and* v_1*,* w_1*,* v_2*,* w_2 *are vectors in* H*, we have*

$$\int_G \langle \pi(g)v_1, w_1\rangle \overline{\langle \pi(g)v_2, w_2\rangle} d\mu(g) = \frac{\langle v_1, v_2\rangle_H \overline{\langle w_2, w_1\rangle}_H}{\dim(H)}. \tag{5.20}$$

Proof. (1) A matrix coefficient of π_1 (resp. π_2) is a vector in the π_1-isotypic component (resp. π_2-isotypic component) of the regular representation of G. These isotypic components are orthogonal if π_1 and π_2 are not isomorphic (Lemma 3.3.15).

(2) Instead of using Schur's Lemma as we did in Chapter 4, we sketch a different argument: we use the fact that the isotypic component $M(\pi) \subset L^2(G)$, under the action of $\mathrm{reg} \boxtimes \mathrm{reg}'$, is isomorphic to $\pi \boxtimes \bar{\pi}$ as a representation of $G \times G$. This is an irreducible unitary representation of $G \times G$, and as such there is on $M(\pi)$ a unique $(G \times G)$-invariant inner product, up to multiplication by a positive scalar. The L^2-inner product, restricted to $M(\pi)$, is such an inner product, but so is the inner product induced by

$$\langle v_1 \otimes w_1, v_2 \otimes w_2\rangle = \langle v_1, v_2\rangle_H \langle w_1, w_2\rangle_{\bar{H}} = \langle v_1, v_2\rangle_H \overline{\langle w_2, w_1\rangle}_H$$

on $H \otimes \bar{H}$. Hence there exists some $\alpha > 0$ such that

$$\int_G \langle \pi(g)v_1, w_1\rangle \overline{\langle \pi(g)v_2, w_2\rangle} d\mu(g) = \alpha \langle v_1, v_2\rangle_H \overline{\langle w_2, w_1\rangle}_H \tag{5.21}$$

for all vectors v_1, w_1, v_2, w_2. In order to compute α, we use the following trick: we fix an arbitrary non-zero vector $v \in H$, and take $v_1 = v_2 = v$ and, successively, $w_1 = w_2 = e_i$, the elements of an orthonormal basis of H. We then sum the identity (5.21) over i, and obtain

$$\sum_i \int_G |\langle \varrho(g)v, e_i\rangle|^2 d\mu(g) = \alpha \|v\|^2 \dim(H).$$

But the left-hand side is equal to

$$\int_G \sum_i |\langle \varrho(g)v, e_i\rangle|^2 d\mu(g) = \int_G \|\pi(g)v\|^2 d\mu(g) = \|v\|^2,$$

using the orthonormality of the basis and the unitarity of $\pi(g)$. Hence, by comparison, we get $\alpha = 1/\dim(H)$, which gives the statement (5.20). □

Proof of Theorem 5.5.1. (1) The space $L^2(G^\sharp)$ is a closed subspace of $L^2(G)$. The characters of finite-dimensional representations are (non-zero) continuous functions, invariant under conjugation, and therefore belong to $L^2(G^\sharp)$. Since the character of $\pi \in \hat{G}$ lies in $M(\pi)$ (as a sum of matrix coefficients), the distinct characters are orthogonal, and Lemma 5.5.2 actually shows that they form an orthonormal system in $L^2(G^\sharp)$.

In order to show its completeness, we need only check that if $\varphi \in L^2(G^\sharp)$ is conjugacy-invariant, its isotypic components, say φ_π, are multiples of the character of π for all irreducible representations π. This can be done by direct computation, just as in the case of finite groups (Section 4.3.3), or by the following argument: it is enough to prove that the space $M(\pi) \cap L^2(G^\sharp)$, in which φ_π lies, is one-dimensional, since we already know that χ_π is a non-zero element of it.

But we can see this space $M(\pi) \cap L^2(G^\sharp)$ as the invariant subspace of $M(\pi)$ when G acts on $L^2(G)$ by the diagonal or conjugation combination of the two regular representations, i.e.,

$$\varrho(g)\varphi(x) = \varphi(x^{-1}gx)$$

for $\varphi \in L^2(G)$. The isotypic component $M(\pi)$ is isomorphic to $H_\pi \otimes \bar{H}_\pi \simeq \mathrm{End}(H_\pi)$ as a vector space, and the corresponding action on $\mathrm{End}(H_\pi)$ is the usual representation of G on an endomorphism space. Thus the G-invariants of $M(\pi)$ under the action ϱ corresponds to the space of G-invariants in $\mathrm{End}(H_\pi)$, which we know is $\mathrm{End}_G(H_\pi) = \mathbf{C}\mathrm{Id}$, by Schur's Lemma. Transporting the identity under the isomorphism of $\mathrm{End}(H_\pi)$ with $M(\pi)$ above leads precisely to the character of π (we leave this as an exercise!), which proves the desired statement.

(2) Because characters are conjugacy-invariant, the action of Φ_π on a unitary representation is an intertwiner (Exercise 5.3.3). In particular, Φ_π acts by multiplication by a scalar on every irreducible unitary representation $\varrho \in \hat{G}$. This scalar is equal to the trace of Φ_π acting on ϱ, divided by $\dim \pi$. Since the trace is given by

$$(\dim \pi) \int_G \overline{\chi_\pi(g)} \chi_\varrho(g) d\mu(g),$$

which is equal to 0 if π is not isomorphic to ϱ, and to $\dim \pi$ otherwise (by orthonormality of characters), we see that Φ_π is the identity on the π-isotypic component of any unitary representation, while it is zero on all other isotypic components. This means that it is the desired projection.

(3) If ϱ is a finite-dimensional unitary representation of G, we have

$$\chi_\varrho = \sum_{\pi \in \hat{G}} n_\pi(\varrho) \chi_\pi,$$

and by orthonormality of the characters we find

$$\langle \chi_\varrho, \chi_\pi \rangle = n_\pi(\varrho)$$

for $\pi \in \hat{G}$. Again orthonormality implies that the formula (5.19) is valid. Applied to $\varrho_1 = \varrho_2 = \varrho$, this gives

$$\|\chi_\varrho\|^2 = \sum_{\pi \in \hat{G}} n_\pi(\varrho)^2.$$

Each term in this sum is a non-negative integer, hence the L^2-norm is equal to 1 if and only if a single term, say $n_\pi(\varrho)$, is non-zero, and in fact equal to 1, which means that ϱ is isomorphic to π. □

Exercise 5.5.3 (Paradox?). Explain why Part (2) of Theorem 5.5.1 does not conflict with the result of Exercise 3.1.5 when G is infinite (note that the action of the projection Φ_π is not given by an element of the group algebra $\mathbf{C}(G)$).

Exercise 5.5.4. Let $\varrho : G \longrightarrow \mathrm{U}(H)$ be a unitary representation of a compact group G, such that, for any $f \in L^2(G)$, the operator $\varrho(f)$ on H is *compact*. Show that the multiplicity of any irreducible representation $\pi \in \hat{G}$ is finite in H.

Remark 5.5.5 (Less duality). All the statements, except for the care needed with L^2-theory, are exactly identical with those which are valid for finite groups. There is, however, at least one sharp difference: there is no good analogue of the second orthogonality formula of Corollary 4.4.1, which expresses the orthogonality of the columns of the character table of a finite group: the expression

$$\sum_{\varrho \in \hat{G}} \chi_\varrho(h) \overline{\chi_\varrho(g)}$$

for g, $h \in G$, does not make sense – in general – in any usual sense (i.e., when G is infinite, this is usually a divergent series.) In other words, the duality between conjugacy classes and irreducible representations is even fuzzier than was the case for finite groups.

Other features of the representations of finite groups that are missing when G is infinite are those properties having to do with integrality properties (though it is tempting to think that maybe there should be some analogue?)

In addition to character theory, matrix coefficients remain available to describe orthonormal bases of $L^2(G,\mu)$. We present this in two forms, one of which is more intrinsic since it does not require a choice of basis. However, it gives an expansion of a different nature than an orthonormal basis in Hilbert space, which is not so well-known in general.

Theorem 5.5.6 (Decomposition of the regular representation). *Let G be a compact topological group and let μ be the probability Haar measure on G.*

(1) *For $\pi\in\hat{G}$, fix an orthonormal basis $(e_{\pi,i})_{1\leqslant i\leqslant \dim(\pi)}$ of the space of π. Then the family of matrix coefficients*

$$\varphi_{\pi,i,j}(g)=\sqrt{\dim(\pi)}\langle\pi(g)e_{\pi,i},e_{\pi,j}\rangle$$

form an orthonormal basis of $L^2(G,\mu)$, where π runs over $\hat{G}$ and $1\leqslant i,j\leqslant\dim\pi$.

(2) *For $\pi\in\hat{G}$, acting on the space H_π, consider the map*

$$A_\pi\begin{cases}L^2(G)\longrightarrow\mathrm{End}(H_\pi)\\ \varphi\mapsto\displaystyle\int_G\varphi(g)\pi(g^{-1})d\mu(g),\end{cases}$$

Then the A_π give "matrix-valued" Fourier coefficients for φ, in the sense that

$$\varphi(x)=\sum_{\pi\in\hat{G}}(\dim\pi)\,\mathrm{Tr}(A_\pi(\varphi)\pi(x)),$$

for all $\varphi\in L^2(G)$, where this series converges in $L^2(G,\mu)$.

Proof. The first statement (1) is a consequence of the Peter-Weyl theorem and the orthogonality of matrix coefficients. The second statement is another formulation of the Peter-Weyl decomposition, because the summands $g\mapsto\mathrm{Tr}(A_\pi(g)\chi_\pi(g))$ are elements of $M(\pi)$. The advantage of (2) is that we obtain an intrinsic decomposition without having to select a basis of the spaces of irreducible representations.

The statement itself is now easy enough: given $\varphi\in L^2(G)$, we have an L^2-convergent series

$$\varphi=\sum_{\pi\in\hat{G}}\varphi_\pi,$$

where φ_π is the orthogonal projection of φ on the π-isotypic component. We compute it using the projection formula of Theorem 5.5.1 applied to the regular representation and to the irreducible representation π. Using the unitarity, this gives

$$\begin{aligned}\varphi_\pi(x)&=(\dim\pi)\int_G\overline{\chi_\pi(g)}\,\mathrm{reg}(g)\varphi(x)d\mu(g)\\&=(\dim\pi)\int_G\chi_\pi(g^{-1})\varphi(xg)d\mu(g)\\&=(\dim\pi)\,\mathrm{Tr}\Big(\int_G\varphi(xg)\pi(g^{-1})d\mu(g)\Big)\\&=(\dim\pi)\,\mathrm{Tr}\Big(\int_G\varphi(y)\pi(y^{-1}x)d\mu(y)\Big)\\&=(\dim\pi)\,\mathrm{Tr}\Big(\Big(\int_G\varphi(y)\pi(y^{-1})d\mu(y)\Big)\pi(x)\Big)=(\dim\pi)\,\mathrm{Tr}(A_\pi(\varphi)\pi(x)),\end{aligned}$$

as claimed. □

Exercise 5.5.7 (G-finite vectors). Let G be a compact group with probability Haar measure μ, and let

$$\varrho : G \longrightarrow \mathrm{U}(H)$$

be a unitary representation of G. A vector $v \in H$ is called *G-finite* if the subrepresentation generated by v is finite-dimensional.

(1) Show that the space H_1 of G-finite vectors is stable under G, and that it is dense in H. When is it a subrepresentation?

(2) Show that a function $f \in L^2(G, \mu)$ is a G-finite vector of the regular representation if and only if f is a finite linear combination of matrix coefficients of irreducible unitary representations of G. (These functions are analogues, for G, of the trigonometric polynomials in the case of the circle group $\mathbf{S}^1$.)

(3) Prove the analogue for a unitary representation ϱ of a compact group G of the property described in Exercise 4.3.29 for finite groups: for any irreducible representation π and any vector v in the π-isotypic component of ϱ, the subrepresentation generated by v is the direct sum of at most $\dim(\pi)$ copies of π.

5.6. Some first examples

We present in this section some simple examples of characters of compact groups.

Example 5.6.1 (Representations of $\mathrm{SU}_2(\mathbf{C})$). The most basic example of non-abelian compact group is probably the group $\mathrm{SU}_2(\mathbf{C}) \subset \mathrm{SL}_2(\mathbf{C})$. We have already seen that it has irreducible representations ϱ_m of degree $m+1$ for all integers $m \geqslant 0$, obtained by restricting the representation of $\mathrm{SL}_2(\mathbf{C})$ on homogeneous polynomials in two variables (see Section 2.6.1 and Exercise 2.7.11).

The concrete incarnation of the Peter-Weyl theorem in that case is the fact that these represent *all* irreducible representations of $\mathrm{SU}_2(\mathbf{C})$. We state this formally:

Theorem 5.6.2 (Irreducible representations of $\mathrm{SU}_2(\mathbf{C})$). *The only irreducible unitary representations of* $\mathrm{SU}_2(\mathbf{C})$ *are given by the representations ϱ_m described above. We have* $\dim \varrho_m = m+1$ *and the character χ_m of ϱ_m is given by*

$$\chi_m\Bigl(\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix}\Bigr) = \frac{\sin((m+1)\theta)}{\sin\theta} \tag{5.22}$$

for $\theta \in [0, \pi]$.

Proof. The definition of ϱ_m makes it clear that it is a continuous representation. Thus we must check that there are no other irreducible unitary representation of $\mathrm{SU}_2(\mathbf{C})$ than those. We will use the completeness criterion from character theory (Theorem 5.5.1) to do this (there are other methods, the most elegant being probably the analysis of representations of the Lie algebra of $\mathrm{SU}_2(\mathbf{C})$).

The set $\mathrm{SU}_2(\mathbf{C})^\sharp$ of conjugacy classes in $\mathrm{SU}_2(\mathbf{C})$ can be identified with the interval $[0, \pi]$ using the map

$$c \left\{ \begin{array}{ccc} [0,\pi] & \longrightarrow & \mathrm{SU}_2(\mathbf{C})^\sharp \\ \theta & \mapsto & \begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix}. \end{array} \right.$$

Indeed, this follows from diagonalizability of unitary matrices in an orthonormal basis, which means that any $g \in \mathrm{SU}_2(\mathbf{C})$ is conjugate in $\mathrm{U}_2(\mathbf{C})$ to such a matrix for some θ, say $g = xc(\theta)x^{-1}$. Replacing x by αx for some $\alpha \in \mathbf{C}$, we can ensure that $\det(x) = 1$, i.e., that $x \in \mathrm{SU}_2(\mathbf{C})$. Next we see if any $c(\theta)$ is conjugate to another; first, conjugating by $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ shows that $c(\theta)$ and $c(-\theta)$ are conjugate, and hence there is always a representative with $\theta \in [0, \pi]$, and there can be no further identification because the trace of $c(\theta)$ is a conjugacy invariant, and $\operatorname{Tr} c(\theta) = 2\cos\theta$, which is an injective function on $[0, \pi]$.

In terms of this identification of $\mathrm{SU}_2(\mathbf{C})^\sharp$, it is not difficult to check that for two (square-integrable) class functions φ_1, φ_2, we have

$$\int_G \varphi_1(g)\overline{\varphi_2(g)}d\mu(g) = \frac{2}{\pi}\int_0^\pi \varphi_1(c(\theta))\overline{\varphi_2(c(\theta))}\sin^2\theta d\theta,$$

and of course we already checked (in (2.46)) that

$$\chi_m(c(\theta)) = \frac{\sin((m+1)\theta)}{\sin\theta}.$$

This means that we are reduced to showing that the functions

$$\varphi_m(\theta) = \frac{\sin((m+1)\theta)}{\sin\theta}, \qquad m \geqslant 0,$$

form an orthonormal basis of the space $H = L^2([0,\pi], \frac{2}{\pi}\sin^2\theta d\theta)$. Because this is close enough to classical Fourier series, we can do this by hand, by reduction to a Fourier expansion, and therefore finish the proof.

Let $\varphi \in H$ be given; we define a function ψ on $[-\pi, \pi]$ by

$$\psi(\theta) = \begin{cases} \varphi(\theta)\sin\theta & \text{if } \theta \geqslant 0, \\ \varphi(-\theta)\sin\theta & \text{if } \theta \leqslant 0, \end{cases}$$

(i.e., we extend by the function $\varphi(\theta)\sin\theta$ to be odd). By definition of H, we see that $\psi \in L^2([-\pi,\pi], d\theta)$. We can therefore expand ψ in Fourier series in this space: we have

$$\psi(\theta) = \sum_{h\in\mathbf{Z}} \alpha_h e^{ih\theta}$$

in $L^2([-\pi,\pi])$, with

$$\alpha_h = \frac{1}{2\pi}\int_{-\pi}^{\pi} \psi(\theta)e^{-ih\theta}d\theta.$$

Since ψ is odd, we have $\alpha_h = -\alpha_{-h}$, and in particular $\alpha_0 = 0$, hence the Fourier expansion on $[0,\pi]$ takes the form

$$\varphi(\theta)\sin\theta = \sum_{h\geqslant 1} \alpha_h 2i\sin(h\theta) = 2i\sum_{m\geqslant 0}\alpha_{m+1}\sin((m+1)\theta),$$

i.e.,

$$\varphi = 2i\sum_{m\geqslant 0}\alpha_{m+1}\varphi_m$$

in $L^2([0,\pi], \frac{2}{\pi}\sin^2\theta d\theta)$. Since we already know that the φ_m are an orthonormal system in this space, it follows that they form an orthonormal basis, as we wanted to prove. □

Example 5.6.3. The proof we gave has at least the advantage that we can easily compute the expansion of various class functions on $\mathrm{SU}_2(\mathbf{C})$, and use results on Fourier series to say something about their convergence with respect to other norms than the L^2-norm (in particular pointwise.).

We can now check the Clebsch-Gordan formula for $\mathrm{SU}_2(\mathbf{C})$ is a single stroke of the pen (compare with the proof of Theorem 2.6.3 that we sketched earlier):

Corollary 5.6.4 (Clebsch-Gordan formula for $\mathrm{SU}_2(\mathbf{C})$). *For all $m \geqslant n \geqslant 0$, the representations ϱ_m of $\mathrm{SU}_2(\mathbf{C})$ satisfy*

$$\varrho_m \otimes \varrho_n \simeq \bigoplus_{0 \leqslant i \leqslant n} \varrho_{m+n-2i}.$$

Proof. One might say it is a matter of checking the identity at the level of characters, i.e., of proving the elementary formula

$$\frac{\sin((m+1)\theta)}{\sin\theta}\frac{\sin((n+1)\theta)}{\sin\theta} = \sum_{0\leqslant i\leqslant n} \frac{\sin((m+n-2i)\theta)}{\sin\theta}$$

for all $\theta \in [0,\pi]$. But more to the point, character theory explains *how to guess* (or find) such a formula: by complete reducibility and orthonormality of characters, we know that

$$\varrho_m \otimes \varrho_n \simeq \bigoplus_{0\leqslant k\leqslant mn-1} \langle \chi_m\chi_n, \chi_k\rangle \varrho_k,$$

(the restriction of the range comes from dimension considerations), and we are therefore reduced to computing the multiplicities

$$\langle \chi_m\chi_n, \chi_k\rangle = \frac{2}{\pi}\int_0^\pi \frac{\sin((m+1)\theta)}{\sin\theta}\frac{\sin((n+1)\theta)}{\sin\theta}\frac{\sin((k+1)\theta)}{\sin\theta}\sin^2(\theta)d\theta.$$

This is, of course, an elementary – if possibly boring – exercise. □

Using the coordinates in (5.7) on $\mathrm{SU}_2(\mathbf{C})$, it is easy to see how the representations ϱ_m, defined as acting on polynomials, can be embedded in the regular representation. Indeed, if we see the coordinates $(a,b) \in \mathbf{C}^2$ of some element $g \in \mathrm{SU}_2(\mathbf{C})$ (which are subject to the condition $|a|^2 + |b|^2 = 1$) as a row vector, a direct matrix multiplication shows that the row vector corresponding to a product gh is the same as the vector-matrix product $(a,b)h$. If we restrict a polynomial $P \in \mathbf{C}[X,Y]$ to $(X,Y) = (a,b)$, this gives an intertwiner from V_m to a space of continuous functions on $\mathrm{SU}_2(\mathbf{C})$. Since this map is non-zero (any basis vector X^iY^{m-i} restricts to a non-zero function), it is an injection

$$V_m \hookrightarrow L^2(\mathrm{SU}_2(\mathbf{C})).$$

According to Vilenkin [**43**], it was first observed by É. Cartan that matrix coefficients or characters of irreducible representations of certain important groups lead to most of the "classical" special functions[8] (of course, the fact that the exponential function is a representation of the additive group of $\mathbf{C}$, or of $\mathbf{R}$ by restriction, is an even older phenomenon that has the same flavor.) We present here some simple instances, related to the group $\mathrm{SU}_2(\mathbf{C})$; more information about these, as well as many additional examples, are found in [**43**].

We begin with the characters of $\mathrm{SU}_2(\mathbf{C})$. As functions of the parameter θ describing the conjugacy class, they are of course completely elementary, but a change of variable adds more subtlety:

Definition 5.6.5 (Chebychev polynomials). For $m \geqslant 0$, there exists a polynomial $P_m \in \mathbf{R}[X]$, of degree m, such that

$$\chi_m(\theta) = P_m(2\cos\theta), \qquad 0 \leqslant \theta \leqslant \pi,$$

[8] Especially the functions that arise in mathematical physics, e.g., Bessel functions.

i.e., such that

$$\frac{\sin((m+1)\theta)}{\sin(\theta)} = P_m(2\cos\theta), \tag{5.23}$$

and the polynomial $U_m = P_m(X/2)$ is called the *m-th Chebychev polynomial of the second kind.*

The fact that the characters of $\mathrm{SU}_2(\mathbf{C})$ form an orthonormal basis of the space of class functions translates into the following fact:

Proposition 5.6.6 (Chebychev polynomials as orthogonal polynomials). *The restrictions to $[-1, 1]$ of the polynomials U_m, $m \geqslant 0$, form an orthonormal basis of the space $L^2([-1, 1], d\nu)$, where ν is the measure supported on $[-1, 1]$ given by*

$$d\nu(t) = \frac{2}{\pi}\sqrt{1-t^2}dt.$$

The justification for this substitution is that the Chebychev polynomials arise in many applications completely independently of any (apparent) consideration of the group $\mathrm{SU}_2(\mathbf{C})$. On the other hand, algebraic properties of the representations ϱ_m can lead to very simple (or very natural) proofs of identities among Chebychev polynomials which might otherwise look quite forbidding if one starts from the definition (5.23), and even more if one begins with an explicit (!) polynomial expansion.

Example 5.6.7. The first few Chebychev polynomials U_m are

$$U_0 = 1, \qquad U_1 = 2X, \qquad U_2 = 4X^2 - 1,$$
$$U_3 = 8X^3 - 4X, \qquad U_4 = 16X^4 - 12X^2 + 1, \qquad \ldots$$

(as one can prove, e.g., by straightforward trigonometric manipulations...)

Exercise 5.6.8 (Playing with Chebychev polynomials). In this exercise, we express the Clebsch-Gordan decomposition of $\varrho_m \otimes \varrho_n$ in terms of expansions of Chebychev polynomials, and deduce some combinatorial identities.

(1) Show that we have

$$U_m(X) = \sum_{0\leqslant j\leqslant m/2} (-1)^j \binom{m-j}{m-2j} (2X)^{m-2j}$$

for all $n \geqslant 0$. [Hint: It is helpful here to interpret the Chebychev polynomials in terms of characters of the larger group $\mathrm{SL}_2(\mathbf{C})$.]

(2) Using the Clebsch-Gordan decomposition, show that for $m \geqslant n \geqslant 0$, we have

$$U_m U_n = \sum_{k=0}^{n} U_{m+n-2k}.$$

(3) Deduce that for $m \geqslant n \geqslant 0$, and $k \leqslant (m+n)/2$, we have

$$\sum_{\substack{i+j=k\\ i\leqslant m/2, j\leqslant n/2}} (-1)^{i+j}\binom{m-i}{m-2i}\binom{n-j}{n-2j} = \sum_{\substack{\ell+t=k\\ \ell\leqslant n}} (-1)^t \binom{m+n-2\ell-t}{m+n-2k}.$$

Example 5.6.9 (Representations of $\mathrm{SO}_3(\mathbf{R})$). The group $\mathrm{SO}_3(\mathbf{R})$ of rotations in a 3-dimensional euclidean space is also very important in applications (we will see it appear prominently in the analysis of the hydrogen atom). As it turns out, it is very closely related to the group $\mathrm{SU}_2(\mathbf{C})$, and this allows us to find easily the representations of $\mathrm{SO}_3(\mathbf{R})$ from those of $\mathrm{SU}_2(\mathbf{C})$.

Proposition 5.6.10. *There exists a continuous surjective group homomorphism*

$$p : \mathrm{SU}_2(\mathbf{C}) \longrightarrow \mathrm{SO}_3(\mathbf{R})$$

such that $\ker p = \{\pm 1\} \subset \mathrm{SU}_2(\mathbf{C})$ *has order* 2*. As a consequence, the irreducible unitary representations of the group* $\mathrm{SO}_3(\mathbf{R})$ *are representations* π_ℓ*,* $\ell \geqslant 0$*, of dimension* $2\ell + 1$*, determined by the condition* $\pi_\ell \circ p = \varrho_{2l}$.

Partial proof. One can write explicitly p in terms of matrices; crossing fingers to avoid typing mistakes, and using the fact that a matrix in $\mathrm{SU}_2(\mathbf{C})$ is uniquely of the form $\begin{pmatrix} z & -\bar{y} \\ y & \bar{z} \end{pmatrix}$ for some y, $z \in \mathbf{C}$ with $|z|^2 + |y|^2 = 1$, this takes the form

$$(5.24)\quad p\Bigl(\begin{pmatrix} a+ib & -c+id \\ c+id & a-ib \end{pmatrix}\Bigr) = \\ \begin{pmatrix} (a^2+b^2)-(c^2+d^2) & 2(bd-ac) & -2(ad+bc) \\ 2(ac+bd) & a^2-b^2-c^2+d^2 & 2(ab-cd) \\ 2(ad-bc) & -2(ab+cd) & a^2-b^2+c^2-d^2 \end{pmatrix}$$

but that is about as enlightening as checking by hand the associativity of the product of matrices of fixed size 4 or more; a good explanation for the existence of p is explained in the next chapter, so we defer a reasonable discussion of this part of the result (see also [**38**, §4.3] for a down-to-earth approach). Note at least that the formula makes it easy to check that $p(g) = 1$ if and only if $g = 1$ or -1.

On the other hand, given the existence of p, we notice that if ϱ is any irreducible unitary representation of $\mathrm{SO}_3(\mathbf{R})$, then $\varrho \circ p$ is an irreducible unitary representation of $\mathrm{SU}_2(\mathbf{C})$. By the previous classification, it is therefore of the form ϱ_m for some $m \geqslant 0$. The question is therefore: for which integers $m \geqslant 0$ is ϱ_m of the form $\varrho \circ p$? The answer is elementary: this happens if and only if $ker(p) \subset \ker(\varrho_m)$, and since $\ker(p)$ has only two elements, this amounts to asking that $\varrho_m(-1) = 1$. The answer can then be obtained from the explicit description of ϱ_m, or by character theory using (5.22): -1 corresponds to $\theta = \pi$ and we have

$$\chi_m(-1) = \lim_{\theta \to \pi} \frac{\sin((m+1)\theta)}{\sin(\theta)} = (-1)^m,$$

so that ϱ_m is obtained from an irreducible representation of $\mathrm{SO}_3(\mathbf{R})$ if and only if $m = 2\ell$ is even. Since different values of $\ell \geqslant 0$ lead to representations of different dimension, this gives the desired correspondence. □

Example 5.6.11 (Infinite products). The following class of examples is quite simple (and does not occur that often in applications), but it is enlightening. By Proposition 2.3.17, we see that the irreducible unitary representations of a direct product $G_1 \times G_2$ of compact groups are of the form

$$\varrho_1 \boxtimes \varrho_2,$$

where ϱ_1 (resp. ϱ_2) ranges over the representations in $\hat{G_1}$ (resp. $\hat{G_2}$). This extends, by induction, to a finite product $G_1 \times \cdots \times G_k$, with irreducible representations

$$\varrho_1 \boxtimes \cdots \boxtimes \varrho_k.$$

We now extend this to infinite products of compact groups. Let I be an arbitrary index set (the interesting case being when I is infinite, say the positive integers), and let $(G_i)_{i\in I}$ be any family of compact topological groups indexed by I (for instance, the family $(\mathrm{GL}_2(\mathbf{F}_p))_p$, where p runs over primes). The product group

$$G = \prod_{i\in I} G_i$$

can be given the product topology. By Tychonov's Theorem, this is a compact topological space. One can check that the product on G is continuous, and hence G is a compact topological group. We now determine its irreducible unitary representations.

There is an abundance of irreducible representations arising from the finite products of the groups: for any finite subset $J \subset I$, we have the projection homomorphism

$$(5.25)\qquad p_J : \begin{cases} G & \longrightarrow & G_J = \prod_{i\in J} G_i \\ (g_i)_{i\in I} & \mapsto & (g_j)_{i\in J} \end{cases}$$

which is continuous (by definition of the product topology), so that any irreducible unitary representation $\boxtimes_{i\in J} \varrho_i$ of the finite product G_J gives by composition an irreducible representation of G, with character

$$\chi((g_i)) = \prod_{i\in J} \chi_{\varrho_i}(g_i).$$

Some of these representations are isomorphic, but this only happens when a component ϱ_i is trivial, in which case we might as well have constructed the character using $G_{J-\{i\}}$ (we leave a formal proof to the reader!) In other words, we have a family of irreducible representations parametrized by a finite – possibly empty – subset J of I, and a family $(\varrho_i)_{i\in J}$ of *non-trivial* irreducible unitary representations of the groups G_i, $i \in J$. In particular, the trivial representation of G arises from $I = \emptyset$, in which case $G_J = 1$.

We now claim that these are the only irreducible unitary representations of G. This statement is easy to prove, but it depends crucially on the topological structure of G. For the proof, we use the completeness criterion from Peter-Weyl theory (Corollary 5.4.8), by proving that the linear span (say V) of the matrix coefficients of those known representations is dense in $L^2(G)$.

For this purpose, it is enough to show that the closure of V contains the continuous functions, since $C(G)$ is itself dense in $L^2(G)$. Let therefore φ be continuous on G. The main point is that the product topology imposes that φ depends "essentially" only on finitely many coordinates. Precisely, let $\varepsilon > 0$ be arbitrary. Since G is compact, the function φ is in fact *uniformly* continuous, in the sense that there exists an open neighborhood U of 1 such that

$$|\varphi(g) - \varphi(h)| \leqslant \varepsilon$$

if $gh^{-1} \in U$.[9] The definition of the product topology shows that, for some finite subset $J \subset I$, the open set U contains a product set

$$V = \{(g_i)_{i\in I} \mid g_i \in V_i \text{ for } i \in J\}$$

for suitable open neighborhoods V_i of 1 in G_i. Intuitively, up to a precision ε, it follows that φ "only depends on the coordinates in J". Let then

$$\varphi_J(g) = \varphi(\tilde{g})$$

where $\tilde{g}_i = g_i$ for $i \in J$ and $\tilde{g}_i = 1$ otherwise. Since $g\tilde{g}^{-1} \in V$, this function on G satisfies

$$|\varphi(g) - \varphi_J(g)| \leqslant \varepsilon$$

[9] We leave the proof as an exercise, adapting the classical case of functions on compact subsets of $\mathbf{R}$.

for all $g \in G$.

But now, φ_J can be identified with a function on G_J. Then, by the Peter-Weyl theorem, we can find a linear combination ψ of matrix coefficients of representations of G_J such that

$$\|\psi - \varphi_J\|_{L^2(G_J)} \leqslant \varepsilon.$$

But it is quite easy to see that the probability Haar measure μ on G is such that its image under the projection p_J is the probability Haar measure μ_J on G_J. This means that when we see ψ (and again φ_J) as functions on G, we still have

$$\|\psi - \varphi_J\|_{L^2(G)} \leqslant \varepsilon.$$

Putting these inequalities together, we obtain

$$\|\varphi - \psi\|_{L^2(G)} \leqslant 2\varepsilon,$$

and as ε was arbitrary, we are done.

In Exercise 6.1.4 in the next chapter, we present a slightly different proof, which directly shows that an irreducible unitary representation of G must be of the "known" type.

CHAPTER 6

Applications of representations of compact groups

This chapter presents some applications of the representation theory of compact groups.

6.1. Compact Lie groups are matrix groups

The first application we present is a rather striking fact of differential geometry: the identification of compact Lie groups with compact subgroups of the linear matrix groups $\mathrm{GL}_n(\mathbf{C})$. We recall the definition of Lie groups first:

Definition 6.1.1 (Lie group). A *Lie group G* is a topological group which is also a topological manifold, i.e., for every $g \in G$, there exists an open neighborhood of g which is homeomorphic to an open ball in some euclidean space $\mathbf{R}^n$, $n \geqslant 0$.

The main result of this section is the fact that a much stronger-looking definition leads to the same class of groups, in the compact case.

Theorem 6.1.2 (Compact Lie groups as matrix groups). *A topological group G is a compact Lie group if and only if there exists some $n \geqslant 1$ such that G is homeomorphic to a closed subgroup of $\mathrm{U}_n(\mathbf{C})$, or to a compact subgroup of $\mathrm{GL}_n(\mathbf{C})$.*

The proof of this result is a very nice combination of basic facts of the theory of Lie groups and of the Peter-Weyl theory. The statement is very powerful: in particular, note that for a closed subgroup of $\mathrm{GL}_n(\mathbf{C})$, the multiplication map is not only continuous (as required by the condition that G be a topological group) but "smooth" in an obvious sense. In fact, as described in a bit more detail in Appendix B, the outcome is that G is not only a topological manifold, but also (in an essentially unique way) a smooth manifold, and even real-analytic manifold, with group operations having the same regularity.

Example 6.1.3. (1) Because of the theorem, it is not surprising (!) that all examples of compact Lie groups that one can write directly are, in fact, obviously compact matrix groups. Such is $\mathrm{U}_n(\mathbf{C})$ are its subgroup $\mathrm{SU}_n(\mathbf{C})$, or the subgroup $\mathbf{T}^n \subset \mathrm{SU}_n(\mathbf{C})$ of diagonal matrices, which can be identified also with the torus $(\mathbf{R}/\mathbf{Z})^n$. One can also consider the group of real orthogonal matrices $\mathrm{O}_n(\mathbf{R})$, which can be seen as $\mathrm{SU}_n(\mathbf{C}) \cap \mathrm{GL}_n(\mathbf{R})$, or the unitary symplectic group $\mathrm{USp}_{2g}(\mathbf{C}) \subset \mathrm{GL}_{2g}(\mathbf{C})$, which consists of unitary matrices of size $2g$ preserving a fixed non-degenerate alternating bilinear form.

(2) One can also give many examples of compact topological groups which are not Lie groups, though it is maybe not immediately obvious that they can not be embedded in a matrix group in any way – this becomes another consequence of the theorem. The infinite products considered in Example 5.6.11, namely

$$G = \prod_{i \in I} G_i, \qquad G_i \neq 1,$$

are of this type, provided I is infinite and infinitely many among the G_i are non-trivial (for the simplest example, take I to be countable and $G_i = \mathbf{Z}/2\mathbf{Z}$ for all i). Indeed, from the description of the irreducible unitary representations of G in Example 5.6.11, we see any one of them contains in its kernel a subgroup

$$G^{(J)} = \prod_{i \notin J} G_i$$

for some finite subset $J \subset I$. A finite-dimensional representation of G, which is a direct sum of finitely many such representations, therefore has kernel containing a finite intersection

$$G^{(J_1)} \cap \cdots \cap G^{(J_k)}$$

of such subgroups, and the latter contains $G^{(J)}$ for $J = J_1 \cup \cdots \cup J_k$, which is a non-trivial group since $I - J$ is not empty. Thus G has no finite-dimensional faithful representation.

(3) Another example is the following group, which is called the group of p-adic integers (it, and its generalizations, are of considerable importance in number theory). Let p be a fixed prime number, and consider first the infinite product

$$G_p = \prod_{k \geqslant 1} \mathbf{Z}/p^k\mathbf{Z}$$

and then its subgroup

$$\mathbf{Z}_p = \{(x_k) \in G_p \mid x_{k+1} \equiv x_k \,(\mathrm{mod}\, p^k)\} \subset G_p.$$

It is again an exercise to check that $\mathbf{Z}_p$ is a closed subgroup of G_p, and hence a compact topological group. It is an abelian group, and one can see as follows that it does not have a faithful finite-dimensional representation. First, since $\mathbf{Z}_p$ is abelian, its irreducible unitary representations are one-dimensional. Then we note that if χ is a character of $\mathbf{Z}_p$, there exists an integer $j \geqslant 1$ such that

$$\ker \chi \supset \{(x_k) \in \mathbf{Z}_p \mid x_j = 0\}$$

(note that if $x_j = 0$, then $x_1 = 0, \ldots, x_{j-1} = 0$, each in its respective group $\mathbf{Z}/p^k\mathbf{Z}$). Indeed, one can see that the sets

$$U_j = \{(x_k) \in \mathbf{Z}_p \mid x_j = 0\}$$

form a fundamental system of neighborhoods of 0 in $\mathbf{Z}_p$. Thus, by continuity, there exists $j \geqslant 0$ such that $\chi(U_j)$ is contained in a fixed neighborhood V of 1 in the circle $\mathbf{S}^1$, for instance the intersection $V = \mathbf{S}^1 \cap \{z \in \mathbf{C} \mid |z-1| < 1/2\}$. But then, since U_j is a subgroup of $\mathbf{Z}_p$, the set $\chi(U_j)$ is a subgroup of $\mathbf{S}^1$ contained in V. But it is elementary that $\{1\}$ is the only such subgroup, which means that $U_j \subset \ker \chi$, as claimed.

Now we use an argument quite similar to that in Example 5.6.11: the kernel of any finite direct sum of characters of $\mathbf{Z}_p$ will also contain a group of the type U_j, and hence such a representation is not faithful.

Exercise 6.1.4 (No small subgroups in unitary groups). (1) Show that in any unitary group $\mathrm{U}_n(\mathbf{C})$, there is a neighborhood V of 1 which contains no non-trivial subgroup.

(2) Use this to reprove the result of Example 5.6.11 by directly showing that any irreducible representation of an infinite product of compact groups is of the "known" form $\varrho \circ p_J$, with notation as in (5.25).

The group $\mathbf{Z}_p$ is abelian, but note that $\mathbf{Z}_p$ is also a topological ring, and one can therefore define groups like

$$\mathrm{SL}_2(\mathbf{Z}_p) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mid a,\ b,\ c,\ d \in \mathbf{Z}_p, \quad ad - bc = 1 \right\}.$$

With the group structure coming from matrix multiplication and the induced topology from the product topology on $\mathbf{Z}_p^4$, this is again a compact topological group (it is a closed subset of the compact space $\mathbf{Z}_p^4$). Of course, it is non-abelian. Its representation theory plays an important role in number theory.

We conclude by mentioning that the similarity between the two counter-examples is not accidental. In fact, a deep theorem of Montgomery-Zippin shows that a topological group G is a Lie group if and only if contains no "small" subgroups, i.e., there is a neighborhood of 1 in G containing no non-trivial subgroup.

We come now to the proof of the theorem. For this, we will need to use without proof the following facts concerning Lie groups:
– A Lie group G has a well-defined *dimension* $\dim(G)$, a non-negative integer, which is the dimension of G as a manifold; for instance

$$\dim(\mathbf{R}) = 1, \qquad \dim(\mathrm{GL}_n(\mathbf{R})) = n^2, \qquad \dim(\mathrm{GL}_n(\mathbf{C})) = 2n^2$$

(the last case illustrates that the dimension involved is that of G as a real manifold.)
– A compact Lie group G has only finitely many connected components; in particular, if G is compact, we have $\dim(G) = 0$ if and only if G is a finite group (which is a compact Lie group with the discrete topology).
– If $H \subset G$ is a closed subgroup of a Lie group, then in fact H is a smooth submanifold, and H is itself a Lie group.
– In the same situation where $H \subset G$ is a closed subgroup, we have in particular $\dim(H) \leqslant \dim(G)$, and if there is equality, the subgroup H is a union of connected components of G; especially, if H and G are both connected, we have $H = G$.

Now we embark on the proof...

Proof of Theorem 6.1.2. In view of the facts above, it is enough to prove the following statement, which introduces representation theory: if G is a compact Lie group, there exists a finite-dimensional *faithful* representation

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of G; indeed, fixing a basis of E, ϱ is then an injective homomorphism of G into $\mathrm{GL}_n(\mathbf{C})$ with $n = \dim(E)$. Since G is compact, ϱ is an homeomorphism onto its image, which is therefore a compact subgroup of $\mathrm{GL}_n(\mathbf{C})$, and with respect to an inner product for which ϱ is unitary, this image is in fact a subgroup of $\mathrm{U}_n(\mathbf{C})$.

The basic idea is now to use the fact that Peter-Weyl theory provides us with many finite-dimensional representations of G, indeed enough to separate points, which means that for every $g \in G$, $g \neq 1$, there is at least one finite-dimensional representation ϱ_g such that $\varrho_g(g) \neq 1$. If g is also in the connected component of 1 in G, then $\ker \varrho_g$ will be a closed subgroup of G with strictly smaller dimension, and we can argue roughly by induction.

We present this slightly differently, merely for the sake of diversity. Let $d \geqslant 0$ be the minimal dimension of the kernel of some finite-dimensional representation ϱ of G. We claim that $d = 0$; if that is the case, and ϱ is such that $\dim \ker \varrho = 0$, we see by the last fact above that the kernel is at most a finite subgroup of G. But then we can also consider the representation

$$\varrho \oplus \bigoplus_{\substack{g \in \ker \varrho \\ g \neq 1}} \varrho_g,$$

which is still finite-dimensional and is now faithful.

Let ϱ be such that the kernel H of ϱ has dimension d. Now assume, for contradiction, that $\dim H = \dim \ker \varrho \geqslant 1$. Then we can find some h which is not trivial, but is in the connected component of H containing 1. Then

$$\ker(\varrho \oplus \varrho_h) = H \cap \ker \varrho_h$$

is a proper subgroup of H. Its dimension is therefore $\leqslant \dim(H)$. But it can not be equal! Indeed, by the facts recalled before the proof, this would only be possible if the connected component of 1 in H coincided with that in $H \cap \ker \varrho_h$, which is not the case as h is in one, but not the other! Thus $\dim \ker(\varrho \oplus \varrho_h) < \dim H = d$, and this is a contradiction with the definition of d, which means the supposition that $d \geqslant 1$ is untenable. □

6.2. The Frobenius-Schur indicator

The results in this section apply equally well (and are of interest!) for finite groups. The basic issue they address is the following: given a finite-dimensional (complex) representation

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

of a compact group G, does there exist on E a symmetric, or alternating, non-degenerate bilinear form b which is invariant under G, i.e., such that

$$b(\varrho(g)v, \varrho(g)w) = b(v, w)$$

for all $g \in G$ and v, $w \in E$? This question should be contrasted with the unitarizability of ϱ, which can be interpreted partly as saying that there is always on E an invariant non-degenerate (positive-definite) hermitian form. As a first illustration of the techniques that will be used, we spell out the following algebraic version of this fact:

Proposition 6.2.1. *Let G be a compact group, $\varrho : G \longrightarrow \mathrm{GL}(E)$ an* irreducible *finite-dimensional complex representation of G. Then there exists, up to multiplication by a non-zero scalar, a unique G-invariant hermitian form b on E. Such a form b is non-degenerate, and in fact b or $-b$ is a positive-definite hermitian form on E making ϱ into a unitary representation.*

Proof. What is new compared with our earlier unitarizability statement is the statement about uniqueness of b up to a constant. □

Remark 6.2.2. It is tempting to consider the real and imaginary parts of an invariant hermitian form to construct symmetric and alternating forms; however, these are only $\mathbf{R}$-bilinear!

Theorem 6.2.3 (Frobenius-Schur indicator). *Let G be a compact group with Haar measure μ, and let*

$$\varrho : G \longrightarrow \mathrm{GL}(E)$$

be a finite-dimensional representation of G. Define the Frobenius-Schur indicator *of ϱ by*

$$\mathrm{FS}(\varrho) = \int_G \chi_\varrho(g^2)\,d\mu(g). \tag{6.1}$$

Then if ϱ is irreducible, we have:

(1) *The representation ϱ is of orthogonal type if and only if* $\mathrm{FS}(\varrho) = 1$*;*

(2) *The representation ϱ is of symplectic type if and only if* $\mathrm{FS}(\varrho) = -1$*;*

(3) *The representation ϱ is of complex type if and only if* $\mathrm{FS}(\varrho) = 0$.

In particular, one, and exactly one, of the three possibilities arise.

Proof. The first part of the argument is relatively similar to that used earlier to "explain" the unitarizability of irreducible representations. Only towards the end do we work out, using character theory, a numerical criterion that distinguishes the three possible types of representations – this naturally introduces the Frobenius-Schur indicator as the "right" tool for this.

We let B denote the vector space of bilinear forms on the space E of ϱ; the group acts on B by the formula

$$(g \cdot b)(v, w) = b(\varrho(g^{-1})v, \varrho(g^{-1})w)$$

for $b \in B$ and v, $w \in E$, and an invariant bilinear form on E is therefore simply an element of the subspace B^G. We must attempt to compute this space, and in particular determine its dimension.

For this, we compute first the character of the action of G on B. This is quite simple: the linear isomorphism

$$\begin{cases} E' \otimes E' \longrightarrow B \\ \lambda_1 \otimes \lambda_2 \mapsto b_{\lambda_1,\lambda_2}, \end{cases}$$

where

$$b_{\lambda_1,\lambda_2}(v, w) = \lambda_1(v)\lambda_2(w) \tag{6.2}$$

is an isomorphism of representations, where E' carries the contragredient of ϱ. Hence by the character formalism, we have

$$\chi_B(g) = \chi_{\tilde{\varrho}}(g)^2 = \overline{\chi_\varrho(g)}^2. \tag{6.3}$$

The next step looks innocuous: by the projection formula on invariants, we have

$$\dim B^G = \int_G \chi_B(g) d\mu(g),$$

and we can bound this from above by

$$\dim B^G \leqslant \int_G |\chi_B(g)| d\mu(g) = \int_G |\chi_\varrho(g)|^2 d\mu(g) = 1, \tag{6.4}$$

so that the space B^G is either zero or one-dimensional, i.e., if there exists a non-zero invariant bilinear form on E, it is unique up to scalar.[1]

What remains to be done is to understand when the dimension is 0 and 1, and this will lead to the refined statement of the theorem. The key is that B has an a-priori decomposition

$$B = B_{sym} \oplus B_{alt} \tag{6.5}$$

into two subrepresentations, where B_{sym} is the space of symmetric bilinear forms and B_{alt} the space of alternating bilinear forms. It is indeed clear that B_{sym} and B_{alt} are G-invariant in B, and the decomposition of B as a vector space is well-known: $B_{sym} \cap B_{alt} = 0$, and one can write any $b \in B$ in the form

$$b = b_s + b_a, \qquad b_s(v, w) = \frac{1}{2}(b(v, w) + b(w, v)), \quad b_a(v, w) = \frac{1}{2}(b(v, w) - b(w, v)), \tag{6.6}$$

[1] As mentioned, this might pass unnoticed, but we have obtained here an upper-bound for the invariants in a representation (the bilinear forms) using information concerning those of a space which seems, a priori, unrelated (the hermitian forms S); this is all done through the remarkable effect of character theory.

with $b_s \in B_{sym}$ and $b_a \in B_{alt}$ (note that (6.5) can be interpreted as the decomposition of B under the representation of the group $\mathfrak{S}_2 = \mathbf{Z}/2\mathbf{Z}$ on B by permutation of the arguments, i.e., the generator $1 \in \mathbf{Z}/2\mathbf{Z}$ acts by $1 \cdot b(v,w) = b(w,v)$).

It follows from (6.5) that

$$B^G = B^G_{sym} \oplus B^G_{alt},$$

with the summands being the spaces of invariant symmetric or alternating bilinear forms. Since $\dim B^G \leqslant 1$, we get the basic trichotomy in terms of bilinear forms: either $\dim B^G_{sym} = 1$, $\dim B^G_{alt} = 0$ (orthogonal type); $\dim B^G_{sym} = 0$, $\dim B^G_{alt} = 1$ (symplectic type); or $B^G_{sym} = B^G_{alt} = B^G = 0$ (complex or unitary type). One might object that (for instance) it is possible that $\dim B^G_{alt} = 1$ but that a non-zero $b \in B^G_{alt}$ is degenerate, whereas the definition of symplectic type asks for a non-degenerate bilinear form. But for any non-zero $b \in B^G$, the kernel of b, i.e., the subspace

$$\ker b = \{v \in E \mid b(v,w) = 0 \text{ for all } w \in E\},$$

is a subrepresentation of E (since for all $v \in \ker b$ and $w \in E$, we have

$$b(\varrho(g)v, w) = b(v, \varrho(g)^{-1}w) = 0$$

so that $\varrho(g)v \in \ker b$). Thus, since E is irreducible and $b \neq 0$ (so that $\ker b \neq E$), we have $\ker b = 0$, and b is non-degenerate.

The numerical criterion for the trichotomy, involving the Frobenius-Schur indicator, arises by noting the following clever way of encapsulating it: the three possibilities are characterized by the value of

$$\dim B^G_{sym} - \dim B^G_{alt} = \begin{cases} 1 - 0 = 1 & \text{for orthogonal type,} \\ 0 - 0 = 0 & \text{for unitary type,} \\ 0 - 1 = -1 & \text{for symplectic type.} \end{cases}$$

which is therefore the "explanation" for the Frobenius-Schur indicator.[2] Again from character theory, we get that the desired invariant is

$$\dim B^G_{sym} - \dim B^G_{alt} = \int_G (\chi_{sym}(g) - \chi_{alt}(g))\, d\mu(g)$$

where, for simplicity, we denote by χ_{sym} and χ_{alt} the characters of B_{sym} and B_{alt}. Therefore we proceed to compute the difference of the two characters.

This is quite easy. For a fixed element $g \in G$, we can diagonalize the unitary operator $\varrho(g)$ in some basis $(e_i)_{1\leqslant i\leqslant n}$ of E, with dual basis (λ_i) of E', so that

$$\varrho(g)e_i = \theta_i e_i, \qquad \tilde{\varrho}(g)\lambda_i = \bar{\theta}_i \lambda_i$$

for some eigenvalues θ_i, whose sum is $\chi_\varrho(g)$ or $\chi_{\tilde{\varrho}}(g)$, respectively. Then the bilinear forms

$$b_{i,j} = b_{\lambda_i,\lambda_j}$$

given by (6.2) form a basis of B, with

$$g \cdot b_{i,j} = \overline{\theta_i\theta_j} b_{i,j},$$

by definition of the action on B. Applying the decomposition (6.6), a basis of B_{sym} is given by the symmetric bilinear forms

$$\frac{1}{2}(b_{i,j} + b_{j,i})$$

[2] Note that we could have exchanged the sign of the two terms, which would just have changed the meaning of the indicators ±1; the choice we made is the standard one, but it is merely a convention.

where i and j are arbitrary (but of course (i,j) and (j,i) give the same basic bilinear form), and a basis of B_{alt} is given by the alternating forms

$$\frac{1}{2}(b_{i,j} - b_{j,i})$$

where this time i and j are arbitrary, but distinct (the pair (i,i) leading to the zero form). Notice that in both case, these are eigenvectors for the action of g with the same eigenvalue $\overline{\theta_i\theta_j}$ (because $b_{i,j}$ and $b_{j,i}$ are in the same eigenspace). Thus we find that

$$\chi_{sym}(g) = \sum_{1\leqslant i\leqslant j\leqslant n} \overline{\theta_i\theta_j}, \qquad \chi_{alt}(g) = \sum_{1\leqslant i<j\leqslant n} \overline{\theta_i\theta_j}.$$

Only the diagonal terms are missing from the second sum compared to the first; hence we get

$$\chi_{sym}(g) - \chi_{alt}(g) = \sum_{1\leqslant i\leqslant n} \overline{\theta_i}^2 = \overline{\chi_\varrho(g^2)}$$

(since the matrix $\varrho(g^2)$ has eigenvalues θ_i^2 in the basis (e_i)), and to conclude and recover the formula (6.1) we may simply observe that

$$\int_G \overline{\chi_\varrho(g^2)}d\mu(g) = \int_G \chi_\varrho(g^2)d\mu(g)$$

since we know already that the integral is a real-number. □

We will now give a few examples. Before doing this, the following result illustrates another interesting meaning of the Frobenius-Schur indicator, and should suggest that complex representations are in the some sense the most usual ones.

Proposition 6.2.4 (Self-dual representations). *Let G be a compact topological group and let ϱ be an irreducible representation of G. Then* $\mathrm{FS}(\varrho) \neq 0$ *if and only if the character of ϱ is real-valued, and this is the case if and only if ϱ is isomorphic to its contragredient representation $\tilde{\varrho}$. Such a representation is called* self-dual.

Proof. According to the proof above, ϱ is symplectic or orthogonal if and only if the space B^G of invariant bilinear forms on the space of ϱ is one-dimensional, and (in view of the character formula (6.3)) this is the case if and only if

$$\int_G \overline{\chi_\varrho(g)}^2 d\mu(g) = 1.$$

As we did earlier, we argue that

$$\Bigl|\int_G \overline{\chi_\varrho(g)}^2 d\mu(g)\Bigr| \leqslant \int_G |\chi_\varrho(g)|^2 d\mu(g) = 1,$$

but now we continue by noticing that *if* there is equality, it must be the case that $\overline{\chi_\varrho(g)}^2$ is proportional to $|\chi_\varrho(g)|^2$, with a scalar multiple of modulus 1. Taking $g=1$ shows that the scalar must be equal to 1, i.e. (taking conjugate), we have

$$\chi_\varrho(g)^2 = |\chi_\varrho(g)|^2 \geqslant 0$$

for all $g \in G$. Since, among complex numbers, only real numbers have a non-negative square, we obtain the first result.

Now the last (and possibly most interesting!) conclusion is easy: since the character of $\tilde{\varrho}$ is $\overline{\chi_\varrho}$, it follows from character theory that ϱ has a real-valued character if and only if it is isomorphic to its contragredient. □

EXAMPLE 6.2.5. (1) Let $G = \mathrm{SL}_2(\mathbf{C})$, or $\mathrm{SU}_2(\mathbf{C})$. Among the representations ϱ_m of G, $m \geqslant 0$, those with m even are of orthogonal type, while those with m odd are of symplectic type. This can be checked in different ways: for $\mathrm{SU}_2(\mathbf{C})$, one may use the integration and character formulas to check that it amounts to proving the identity

$$\frac{2}{\pi}\int_0^{\pi} \frac{\sin(2(m+1)\theta)}{\sin 2\theta} \sin^2\theta d\theta = (-1)^m.$$

For either group, one may also define explicitly an invariant non-degenerate bilinear b form on the space V_m of homogeneous polynomials of degree m in two variables, by putting

$$b(e_i, e_j) = \frac{(-1)^i\delta(i, m-j)}{\binom{m}{i}}$$

for the basis vectors $e_i = X^iY^{m-i}$.

Such a definition certainly defines a bilinear form on V_m, and it is symmetric for m even, alternating for m odd (where $\delta(i, m-i)$ is always zero). To see that it is non-degenerate, observe that the non-zero coefficients of the matrix $(b(e_i, e_j))_{i,j}$ are exactly the anti-diagonal ones, so that the determinant is their product, up to sign, which is non-zero.

It is not immediately obvious, on the other hand, that b is invariant. A fairly quick algebraic proof of this is explained in [**39**, 3.1.4, 3.1.5]: the group $\mathrm{SL}_2(\mathbf{C})$ is generated by the elements

$$u(t) = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}, \quad t \in \mathbf{C}, \qquad a(x) = \begin{pmatrix} x & 0 \\ 0 & x^{-1} \end{pmatrix}, \quad x \in \mathbf{C}^\times, \qquad w = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

so that it is enough to check that

$$b(\varrho_m(g)e_i, \varrho_m(g)e_j) = b(e_i, e_j)$$

for g in one of these three classes (and all basis vectors). We leave the easy cases of $a(x)$ and w to the reader, and just present the (maybe somewhat mysterious) case of $g = u(t)$. In that case, we have

$$\varrho_m(g)e_i = X^i(tX+Y)^{m-i} = \sum_{k=0}^{m-i}\binom{m-i}{k} t^k e_{i+k} \tag{6.7}$$

by the binomial theorem, and hence

$$b(\varrho_m(g)e_i, \varrho_m(g)e_j) = \sum_{k=0}^{m-i}\sum_{\ell=0}^{m-j} t^{k+\ell}\binom{m-i}{k}\binom{m-j}{\ell} b(e_{i+k}, e_{j+\ell}).$$

Using the definition of b, only terms with $i+k+j+\ell = m$ remain, and this can be rearranged as

$$b(\varrho_m(g)e_i, \varrho_m(g)e_j) = (-1)^i t^{m-i-j} \sum_{k=0}^{m-i-j} (-1)^k \frac{\binom{m-i}{k}\binom{m-j}{m-i-j-k}}{\binom{m}{i+k}}$$

(where it is possible that the sum be empty). If one rearranges the ratio of binomial coefficients in terms of factorials, this becomes

$$b(\varrho_m(g)e_i, \varrho_m(g)e_j) = \frac{(-1)^i t^{m-i-j}}{(m-i-j)!}\frac{(m-i)!(m-j)!}{m!}\sum_{k=0}^{m-i-j}(-1)^k\binom{m-i-j}{k}.$$

Now the inner sum over k is zero (hence equal to $b(e_i, e_j)$) except when $m = i + j$, and in that last case we also get

$$b(\varrho_m(g)e_i, \varrho_m(g)e_j) = \frac{(-1)^i (m-i)! i!}{m!} = b(e_i, e_j).$$

This example can be considered as a rather striking illustration of the power of character theory: the existence of a symmetric or alternating invariant bilinear form on the space of ϱ_m is obtained by a simple integral of trigonometric functions, but the actual bilinear form is quite intricate and it is not straightforward at all to guess its expression. In particular, note that it is quite different from the $\mathrm{SU}_2(\mathbf{C})$-invariant inner product, for which the vectors e_i are orthogonal (see Example 5.2.12). In fact, this inner product $\langle\cdot,\cdot\rangle$ on V_m is *not* invariant under the larger group $\mathrm{SL}_2(\mathbf{C})$, since ϱ_m is not unitary as a representation of $\mathrm{SL}_2(\mathbf{C})$.

Concretely, recall that the inner product can be defined (a scalar multiple of the one computed in Example 5.2.12) so that the (e_i) are orthogonal and have squared length

$$\langle e_i, e_i \rangle = \frac{1}{\binom{m}{i}}.$$

Then one sees for instance, from (6.7), that

$$\langle \varrho_m(u(t))e_i, \varrho_m(u(t))e_i \rangle = \sum_{k=0}^{m-i} |t|^{2k} \binom{m-i}{k}^2 \binom{m}{i+k}^{-1}$$

for $t \in \mathbf{C}$ and $0 \leqslant i \leqslant m$. This is obviously a non-constant function of t.

(2) This last remark illustrates again the strength of character theory: let us write the crucial inequality (6.4) in the form

$$\dim B^G \leqslant \dim S^G,$$

(with B the space of bilinear forms, S the space of hermitian forms). Now although both sides are purely algebraic invariants that may be defined for any finite-dimensional complex representations of any group, and although the inequality is valid for all compact groups, *it is not* universally valid for finite-dimensional representations of topological groups! Indeed, already for $G = \mathrm{SL}_2(\mathbf{C})$ and $\varrho = \varrho_m$, the right-hand side is 0, while the left-hand side is always 1 (since we checked that the bilinear form b above was $\mathrm{SL}_2(\mathbf{C})$-invariant, and not merely in $B^{\mathrm{SU}_2(\mathbf{C})}$.)

(3) But there is even more to this story: if we write (6.4) in the form

$$\dim B^G \leqslant 1,$$

then it turns out that it *is* valid for any finite-dimensional irreducible representation of any group G (even without imposing continuity conditions)! This is due to the "algebraic" nature of the condition that a bilinear form be invariant, as we will explain in Section 7.1.

(4) For some important classes of finite groups, all representations are of orthogonal type. This applies, for instance, to the symmetric groups (because they can be constructed as matrix representations with values in $\mathrm{GL}_n(\mathbf{Q})$). An interesting consequence of this arises as the application of the following simple lemma:

Lemma 6.2.6 (Groups with all representations orthogonal). *Let G be a finite group such that* $\mathrm{FS}(\varrho) = 1$ *for all irreducible complex representations ϱ of G. Then the sum*

$$\sum_{\varrho \in \hat{G}} \dim(\varrho)$$

of the dimensions of irreducible representations of G is equal to the number of elements of order 2 in G.

Proof. This is quite a cute argument: by assumption, we have

$$\frac{1}{|G|}\sum_{g\in G}\chi_\varrho(g^2)=1$$

for all $\varrho\in\hat{G}$. Multiplying by $\dim(\varrho)$ and then summing over all ϱ, we obtain

$$\begin{aligned}\sum_{\varrho\in\hat{G}}\dim(\varrho)&=\frac{1}{|G|}\sum_{\varrho\in\hat{G}}\sum_{g\in G}\chi_\varrho(g^2)\dim(\varrho)\\&=\frac{1}{|G|}\sum_{g\in G}\sum_{\varrho\in\hat{G}}\chi_\varrho(g^2)\chi_\varrho(1).\end{aligned}$$

By the second orthogonality relation (4.28), the inner sum vanishes unless g^2 is conjugate to 1, i.e., unless $g^2=1$, and in that case it is equal to $|G|$. Thus we get

$$\sum_{\varrho\in\hat{G}}\dim(\varrho)=|\{g\in G\ \mid\ g^2=1\}|,$$

as claimed. □

The question of evaluating this sum was mentioned briefly in Remark 4.2.5.

Exercise 6.2.7. Consider the examples of finite groups for which we computed the full character table (in Section 4.6.2, 4.6.3 and 4.6.4), and for each of them determine the Frobenius-Schur indicators (in particular determine which are self-dual). [Hint: For $\mathrm{GL}_2(\mathbf{F}_p)$, one can use Exercise 4.6.15 to first find very easily the self-dual representations.]

6.3. The Larsen alternative

Our next application has some common features with the Frobenius-Schur theory, but it is a much more recent development which is really a fact about compact, infinite, Lie groups. The results are due to M. Larsen [**27**, §3], and have been extensively developed by N. Katz (for instance in [**22**]).

Their basic motivation can be described as follows: a compact group $G\subset \mathrm{U}_n(\mathbf{C})$ is given, by some means or other, and the question that arises is to identify it, in particular, to prove that it is "big" in some sense. Here, "big" has roughly the following meaning: either one would like to prove that $G\supset \mathrm{SU}_n(\mathbf{C})$, or one knows – again, one way or another – that G preserves either a symmetric or alternating non-degenerate bilinear form, and the goal is to prove that G contains either the corresponding (real) special orthogonal group or the unitary symplectic group. For this, Larsen found a beautiful numerical criterion. We present it here as an interesting and relatively elementary fact about representations of compact groups. It might not be clear whether this is actually applicable in practice, but we will describe quickly in a remark how the problem appears in concrete applications, and how it has been applied by Katz in particular.

The invariant introduced by Larsen is the following:

Definition 6.3.1 (Fourth moment of a representation). Let G be a compact subgroup of $\mathrm{U}_n(\mathbf{C})$ for some $n\geqslant 1$ with Haar measure μ. The *fourth moment* of G is defined by

$$\mathrm{M}_4(G)=\int_G|\operatorname{Tr}(g)|^4d\mu(g).\tag{6.8}$$

More generally, given a finite-dimensional representation ϱ of G, the *fourth moment of ϱ* is defined by

$$\mathrm{M}_4(\varrho) = \int_G |\chi_\varrho(g)|^4 d\mu(g).$$

Thus $\mathrm{M}_4(G)$ corresponds to taking as ϱ the given "tautological" faithful representation $\varrho : G \hookrightarrow \mathrm{U}_n(\mathbf{C})$.

A priori, this would be an arbitrary non-negative real number. However, as in the case of the Frobenius-Schur indicator (6.1), it is in fact an integer, and certain of its values carry important meaning. More precisely, we have the following rather remarkable result of Larsen:

Theorem 6.3.2 (Larsen alternative for unitary groups). *Let $n \geqslant 2$, $G \subset \mathrm{SU}_n(\mathbf{C})$ a compact group. If the fourth moment $\mathrm{M}_4(G)$ is equal to 2, then either G is* finite*, or $G = \mathrm{SU}_n(\mathbf{C})$. In particular, if G is connected, we have $G = \mathrm{SU}_n(\mathbf{C})$.*

The proof is a very nice application of basic character theory and representation theory, together with some basic facts of Lie theory. The first step, which we take "backwards" in comparison with Section 6.2, is to interpret the fourth moment in purely algebraic terms.

Lemma 6.3.3. *Let G be a compact group and*

$$\varrho : G \longrightarrow \mathrm{GL}(V)$$

a finite-dimensional representation of G.

(1) *We have*

$$\mathrm{M}_4(\varrho) = \dim(\mathrm{End}(\varrho) \otimes \mathrm{End}(\varrho))^G = \dim \mathrm{End}(\varrho \otimes \tilde{\varrho})^G. \tag{6.9}$$

(2) *Let π be any of the representations of G on $\varrho \otimes \varrho$, $\varrho \otimes \tilde{\varrho}$ or* $\mathrm{End}(\varrho)$*. If we have a decomposition*

$$\varrho \simeq \bigoplus_i n_i \varrho_i, \qquad n_i \geqslant 0,$$

into G-stable subspaces, with non necessarily irreducible subrepresentations ϱ_i, then we have

$$\mathrm{M}_4(\varrho) \geqslant \sum_i n_i^2,$$

with equality if and only if the ϱ_i are pairwise distinct irreducible representations.

(3) *If $G \subset H$ are compact subgroups of $\mathrm{U}_n(\mathbf{C})$, then we have*

$$\mathrm{M}_4(H) \leqslant \mathrm{M}_4(G). \tag{6.10}$$

Note that (6.9) provides a definition of $\mathrm{M}_4(\varrho)$ which could be used for any (finite-dimensional) representation of any group.

Proof. Note that the fourth moment is an inner product

$$\mathrm{M}_4(\varrho) = \langle |\chi_\varrho|^4, 1\rangle.$$

By the formalism of characters, the function $|\chi_\varrho|^4$ is the character of the representation

$$\tau = \varrho \otimes \varrho \otimes \tilde{\varrho} \otimes \tilde{\varrho},$$

so that $\mathrm{M}_4(\varrho)$ is the dimension of the invariant space τ^G. But using the associativity of the tensor product, and the relations

$$\widetilde{\varrho_1 \otimes \varrho_2} = \tilde{\varrho}_1 \otimes \tilde{\varrho}_2, \qquad \tilde{\tilde{\varrho}} = \varrho,$$

we can arrange the tensor product τ in two ways: either

$$\tau = (\varrho \otimes \varrho) \otimes (\widetilde{\varrho \otimes \varrho}) \simeq \mathrm{End}(\varrho \otimes \varrho),$$

which gives

$$\mathrm{M}_4(\varrho) = \dim(\mathrm{End}(\varrho \otimes \varrho))^G,$$

or

$$\tau = (\varrho \otimes \tilde{\varrho}) \otimes \widetilde{\varrho \otimes \tilde{\varrho}} \simeq \mathrm{End}(\varrho \otimes \tilde{\varrho}) = \mathrm{End}(\mathrm{End}(\varrho)),$$

so that

$$\mathrm{M}_4(\varrho) = \dim(\mathrm{End}(\mathrm{End}(\varrho))^G.$$

This proves (1), and (2) is a general fact about $\dim \mathrm{End}(\pi)^G$ for any representation π: we have

$$\langle \mathrm{End}(\pi), 1\rangle = \sum_{i,j} n_i n_j \langle \varrho_i, \varrho_j\rangle$$

by linearity. Each term is a non-negative integer, and hence

$$\langle \mathrm{End}(\pi), 1\rangle \geqslant \sum_i n_i^2 \langle \varrho_i, \varrho_i\rangle \geqslant \sum_i n_i^2,$$

by keeping only the diagonal terms $i = j$. If there is equality, we see that we must have $\langle \varrho_i, \varrho_j\rangle = \delta(i, j)$, which means that the ϱ_i are irreducible (taking $i = j$) and distinct (for $i \neq j$).

Finally the inequality (6.10), though not at all obvious from the definition (6.8), is clear from (1): if $G \subset H$, then – for any representation of H – the space of G-invariants contains the space of H-invariants. □

Proof of the Larsen alternative. We begin y assuming that G is *not* finite. Thus, we must show that $G = \mathrm{SU}_n(\mathbf{C})$ if and only if $\mathrm{M}_4(G) = 2$.

To approach $\mathrm{M}_4(G)$, we use (2) for the representation of G on the linear space $\mathrm{End}(\mathbf{C}^n)$, i.e., on $\mathrm{End}(\varrho)$ in terms of the defining representation

$$\varrho \,:\, G \hookrightarrow \mathrm{U}_n(\mathbf{C}).$$

We recall that this representation is the *conjugation* action, i.e., that

$$g \cdot A = gAg^{-1}$$

for $g \in G$ and $A \in E = \mathrm{End}(\mathbf{C}^n)$ (it is the restriction of the corresponding action for $\mathrm{SU}_n(\mathbf{C})$, or indeed for $\mathrm{GL}_n(\mathbf{C})$). There is, as usual, a canonical invariant subspace of dimension one, namely $\mathbf{C}\mathrm{Id} \subset E$. Moreover, a stable (orthogonal) complement is

$$E_0 = \{A \in V \mid \mathrm{Tr}(A) = 0\},$$

the space of endomorphisms of trace 0. Hence we have a first decomposition into subrepresentations

$$E = \mathbf{C}\mathrm{Id} \oplus E_0. \tag{6.11}$$

If only for dimension reasons, the two components are non-isomorphic; therefore, by (2) in the previous lemma, we get automatically

$$\mathrm{M}_4(G) \geqslant 1^2 + 1^2 = 2.$$

Thus, we see first that $\mathrm{M}_4(\mathrm{SU}_n(\mathbf{C})) = 2$ means that the decomposition (6.11) is a decomposition into irreducible representations in the case of $\mathrm{SU}_n(\mathbf{C})$; this is indeed the case: the first component, because it is one-dimensional, and the second by Exercise 2.7.12.

By the same token, we also see that if $G \subset \mathrm{SU}_n(\mathbf{C})$, we can only have $\mathrm{M}_4(G) = 2$ if E_0 is also irreducible as a representation of G. Thus we have to find a non-trivial

G-subrepresentation of E_0. To do this, we must appeal to the fact that G is a Lie group. Thus we consider the tangent space of G at the identity element, which is its Lie algebra,[3] denoted $\mathrm{Lie}(G)$. This is a *real* vector space, of dimension equal to the dimension of G as a manifold. The point is that G acts linearly on $\mathrm{Lie}(G)$, by means of the so-called Adjoint representation, which is obtained by differentiating at the identity the conjugation action of G on itself: denoting by $i(g)$ the inner automorphism that maps x to gxg^{-1}, the adjoint representation is given by

$$\mathrm{Ad}\ \begin{cases} G & \longrightarrow & \mathrm{GL}(\mathrm{Lie}(G)) \\ g & \mapsto & d_1(i(g)) \end{cases}$$

(indeed, this is a well-defined linear map on $\mathrm{Lie}(G)$ since $i(g)(1) = 1$, and it is a representation because $i(gh) = i(g)i(h)$ and $i(g^{-1}) = i(g)^{-1}$.)

This is a *real* representation, since $\mathrm{Lie}(G)$ is a real vector space. Most crucial for us, it has the following property, which is almost immediate: if $G \subset H$, with H also a compact Lie group, then $\mathrm{Lie}(G) \subset \mathrm{Lie}(H)$, and the adjoint representation of G is the restriction of the adjoint representation of H. Applied to $G \subset \mathrm{SU}_n(\mathbf{C})$, it follows that $\mathrm{Lie}(G)$ is a subrepresentation of $\mathrm{Lie}(\mathrm{SU}_n(\mathbf{C}))$.

This is the source of the desired subrepresentation of V^0. In fact, we will now check the following facts:

– The Lie algebra L_n of $\mathrm{SU}_n(\mathbf{C})$ is a real subspace of V^0, such that $L_n \oplus iL_n = L_n \otimes \mathbf{C} = V^0$;

– In fact the adjoint representation on L_n is a real subrepresentation of V^0, i.e., on $L_n \subset V^0$, the adjoint representation is given by $g \cdot A = gAg^{-1}$ for $A \in L_n \subset V^0$.

If we assume these facts, we are done: indeed, for $G \subset \mathrm{SU}_n(\mathbf{C})$, it follows that $\mathrm{Lie}(G) \otimes \mathbf{C} \subset L_n \otimes \mathbf{C} = V^0$ is a subrepresentation. Since we have already claimed – with proof to come! – that V^0 is irreducible as a representation of $\mathrm{SU}_n(\mathbf{C})$, this is only possible if either $\mathrm{Lie}(G)$ is 0 – which means that G is finite – or if $\mathrm{Lie}(G)$ is equal to L_n. In that case, by Lie theory, we have $G = \mathrm{SU}_n(\mathbf{C})$, and therefore the Larsen alternative is proved.

Now we explain the facts mentioned above – these are quite standard, and the reader may well have already encountered them. To begin with, the special unitary group is defined by the conditions

$$\det(g) = 1, \qquad gg^* = 1$$

in $\mathrm{GL}_n(\mathbf{C})$. The tangent space at 1 is obtained by considering the linearized forms of these equations, viewed as applying to matrices A in $\mathrm{M}_n(\mathbf{C})$, which form the tangent space at 1 of $\mathrm{GL}_n(\mathbf{C})$. The first equation becomes $\mathrm{Tr}(A) = 0$, which means $A \in V^0$, and the second becomes

$$A + A^* = 0,$$

i.e., A is skew-hermitian, so

$$L_n = \{A \in \mathrm{M}_n(\mathbf{C}) \mid A = -A^*, \qquad \mathrm{Tr}(A) = 0\} \subset V^0 \tag{6.12}$$

(note that since the adjoint operation $A \mapsto A^*$ is not complex-linear, this is indeed only a real vector space.)

We can easily check explicitly that $V^0 = L_n \otimes \mathbf{C}$: for $A \in V^0$, we write

$$A = \frac{A + A^*}{2} + \frac{A - A^*}{2} = iB + C, \quad \text{(say.)}$$

[3] We will not need the structure of Lie algebra that exists on this space.

Then $C^* = -C$, so C is skew-hermitian, and $B = (A+A^*)/(2i)$ has also $B^* = -(A^*+A)/(2i) = -B$, so that B is skew-hermitian. Since $\mathrm{Tr}(B) = \mathrm{Re}(\mathrm{Tr}(A))$ and $\mathrm{Tr}(C) = i\,\mathrm{Im}(\mathrm{Tr}(A))$, we also have $\mathrm{Tr}(B) = \mathrm{Tr}(C) = 0$, so that we have found a decomposition of A as $C + iB$ with C, B both in L_n. This decomposition is unique, because $L_n \cap iL_n = 0$ (in V^0): any matrix in the intersection is both hermitian and skew-hermitian. So this proves the first claim.

The second one is not too surprising since the adjoint representation is defined using conjugation. To be precise, let $A \in L_n$ be a tangent vector; that elementary differential geometry tells us that $\mathrm{Ad}(g)(A)$ can be computed as

$$\frac{d}{dt} i(g)(x_t)\Big|_{t=0}$$

where $x_t \in \mathrm{SU}_n(\mathbf{C})$ defines any smooth curve with tangent vector A at $t = 0$. As usual, one takes $x_t = \exp(tA)$, where the exponential is that of matrices; then we have

$$i(g)x_t = g\exp(tA)g^{-1} = \exp(tgAg^{-1}),$$

(e.g., using the Taylor series expansion) and the derivative at $t = 0$ gives $\mathrm{Ad}(g)A = gAg^{-1}$, as desired. □

Remark 6.3.4 (The Larsen alternative for other groups). In addition to the case of the unitary group considered above, there are criteria for orthogonal and symplectic groups.

Remark 6.3.5 (Finite groups with $\mathrm{M}_4 = 2$). As observed by Katz [**22**, 1.6.1], there do exist finite groups $G \subset \mathrm{SU}_n(\mathbf{C})$, for some $n \geqslant 2$, for which $\mathrm{M}_4(G) = 2$. For instance, let $G = \mathrm{PSL}_2(\mathbf{F}_7)$; it follows from the character table of $\mathrm{SL}_2(\mathbf{F}_7)$ that G has two distinct irreducible representations π_1 and π_2 of dimension $3 = (7-1)/2$. Unitarized, either of these gives a homomorphism

$$G \longrightarrow \mathrm{U}_3(\mathbf{C}).$$

Since G is a simple group, this is necessarily a faithful representation, and (for the same reason) the composite $G \hookrightarrow \mathrm{U}_3(\mathbf{C}) \xrightarrow{\det} \mathbf{C}^\times$, which can not be injective, is trivial. Thus the image of either of these representations is a finite subgroup of $\mathrm{U}_3(\mathbf{C})$, and one can check that these have fourth moment equal to 2, i.e., that $\mathrm{M}_4(\pi_1) = \mathrm{M}_4(\pi_2) = 2$.

Remark 6.3.6 (How does one apply the Larsen alternative?). We explain here, with a specific example, some of the situations where results like the Larsen alternative are very valuable tools. As already hinted, sometimes theory gives the existence of some group which carries information concerning objects of interest. A very good example, though it is not directly relevant to the Larsen alternative, is the Galois group of the splitting field of a polynomial. This is a finite group, constructed abstractly. If one knows the coefficients of the polynomial, however, it is not so easy to determine the Galois group. In fact, often the only obvious information is that it is isomorphic to a subgroup of $\mathfrak{S}_n$, where n is the degree of the polynomial (for instance, can you guess the Galois group of the splitting field of

$$X^8 - 4X^7 + 8X^6 - 11X^5 + 12X^4 - 10X^3 + 6X^2 - 3X + 2$$

over $\mathbf{Q}$?) In fact, part of what makes the Larsen alternative surprising is that it does not really have an analogue for Galois groups!

Now for the example, which is based on very recent (and very deep) work of N. Katz [**23**]. Fix an integer $d \geqslant 1$ and a prime number p such that $p \nmid d(d-1)$. For any finite field $\mathbf{F}_q$ of order q which is a power of p, and any non-trivial (one-dimensional) character χ of the multiplicative group $\mathbf{F}_q^\times$, one defines the sum

$$S(\chi; q) = \sum_{x \in \mathbf{F}_q} \chi(x^d - dx - 1) \tag{6.13}$$

using the convention $\chi(0) = 0$. These are apparently just complex numbers, but they turn out to be related to some compact Lie groups. Indeed, it follows from the work of A. Weil[4] that for every such pair (q, χ), there exists a well-defined *conjugacy class* $\theta(\chi; q)$ in the unitary group $\mathrm{U}_{d-1}(\mathbf{C})$ such that

$$\operatorname{Tr} \theta(\chi; q) = -\frac{S(\chi; q)}{\sqrt{q}} \tag{6.14}$$

(in particular, note that this implies that

$$|S(\chi; q)| \leqslant (d-1)\sqrt{q},$$

which the reader may try to prove directly, knowing that, even for the first difficult case $d = 3$, all known proofs are very involved...)

The connection with the Larsen alternative arises from the following fact, which is a recent theorem of Katz (closely related to another very deep result of Deligne): there exists a compact subgroup $K \subset \mathrm{U}_{d-1}(\mathbf{C})$, depending a priori on p and d, such that, first, all $\theta(\chi; q)$ are in fact naturally conjugacy classes of K, and second, they become *equidistributed* among conjugacy classes of K, in the sense that for any conjugacy-invariant continuous function $f : K \longrightarrow \mathbf{C}$, we have

$$\int_K f(x) d\mu(x) = \lim_{q \to +\infty} \frac{1}{q-2} \sum_{\chi \neq 1} f(\theta(\chi; q)). \tag{6.15}$$

where μ is the probability Haar measure on K and the sum is over non-trivial characters of $\mathbf{F}_q^\times$.

Thus, if one succeeds in determining what the group K is – something which, just as was the case for Galois group, is by no means clear by just looking at the sums (6.13)! – one can answer many questions about the asymptotic distribution of the sum, something which is of great interest (at least, to arithmeticians...)

Now it is clear why the Larsen alternative is useful: applying first (6.15) with $f(x) = |\operatorname{Tr}(x)|^4$ and then (6.14), we get the alternative formula

$$\begin{aligned} \mathrm{M}_4(K) &= \lim_{q \to +\infty} \frac{1}{q-2} \sum_{\chi \neq 1} |\operatorname{Tr} \theta(\chi; q)|^4 \\ &= \lim_{q \to +\infty} \frac{1}{q^2(q-2)} \sum_{\chi \neq 1} \Big| \sum_{x \in \mathbf{F}_q} \chi(x^d - dx - 1) | \Big|^4 \end{aligned}$$

for the fourth moment of K, which involves the given, concrete, data defining the problem. We may have a chance to evaluate this...

As it turns out, one can show that the compact group K, if $d > 6$ at least, does satisfy $\mathrm{M}_4(K) = 2$ (though proving this is actually quite difficult). Hence the Larsen alternative shows that either K is finite, or $K \supset \mathrm{SU}_{d-1}(\mathbf{C})$. One can analyze further the situation, and the conclusion (still for $d > 6$) is that K is equal to the full unitary group U_{d-1}.

[4] This is a special case of the *Riemann Hypothesis for curves over finite fields.*

In the works of Katz, many other (more general) situations are considered. Note that, even though the statements can be rather concrete, there is no known elementary proof of the deep connection between sums like $S(\chi; p, d)$ and a compact Lie group.

Exercise 6.3.7 (Other moments). One can define other types of moments. For instance, given a compact group G and a finite-dimensional unitary representation ϱ of G, let

$$M_a(\varrho) = \int_G \chi_\varrho(g)^a d\mu(g)$$

for an integer $a \geqslant 0$. It is an elementary consequence of character theory, which is not necessarily clear at first when expressed for a "concrete" group, that $M_a(\varrho)$ is a non-negative integer, as the multiplicity of the trivial representation in the finite-dimensional representation $\varrho^{\otimes a}$.

The sequence of moments $(M_a(\varrho))_{a\geqslant 0}$, as $a \geqslant 0$ varies, can be quite interesting...

(1) Take $G = \mathrm{SU}_2(\mathbf{C})$ and ϱ the tautological inclusion $\mathrm{SU}_2(\mathbf{C}) \hookrightarrow \mathrm{GL}_2(\mathbf{C})$. Show that

$$M_a(\varrho) = 0$$

if a is odd and

$$M_{2a}(\varrho) = \frac{1}{2a+1}\binom{2a}{a}$$

for $a \geqslant 0$. Can you prove directly that the right-hand side is an integer?

(2) Compute the first few terms and identify this sequence in the "Online Encyclopedia of Integer Sequences" (`http://oeis.org`).

(3) For a prime number p, and an element $\alpha \in \mathbf{F}_p^\times$, let

$$S(\alpha; p) = \frac{1}{\sqrt{p}} \sum_{x\in\mathbf{F}_p^\times} e\Big(\frac{x + \alpha\bar{x}}{p}\Big),$$

where $e(\cdot)$ is the character $e(z) = e^{2i\pi z}$ of $\mathbf{R}/\mathbf{Z}$ and $\bar{x}$ designates the inverse of x modulo p (i.e., $x\bar{x} = 1 \,(\mathrm{mod}\, p)$). For reasonably large values of p (say $p \leqslant 100000$) and the first few $a \geqslant 0$, compute (using a computer) the "empirical" moments

$$m_{a,p} = \frac{1}{p-1} \sum_{\alpha\in\mathbf{F}_p^\times} S(\alpha, p)^a.$$

Discuss the behavior of the result as p grows...

Remark 6.3.8 (From $\mathrm{SU}_2(\mathbf{C})$ to $\mathrm{SO}_3(\mathbf{R})$). The Adjoint representation turns out to provide the conceptual explanation of the projection homomorphism

$$p : \mathrm{SU}_2(\mathbf{C}) \longrightarrow \mathrm{SO}_3(\mathbf{R})$$

of Proposition 5.6.10. Indeed, for the compact Lie group $G = \mathrm{SU}_2(\mathbf{C})$, the Lie algebra is a three-dimensional real vector space (by (6.12): $Mat_2(\mathbf{C})$ has dimension 8, the skew-hermitian condition implies that the bottom-left coefficient is minus the conjugate of the top-right one, and that the diagonal ones are purely imaginary, leaving $8 - 2 - 2 = 4$ dimensions, and the matrices of trace zero form a 3-dimensional subspace. In fact,

$$L_2 = \Big\{ \begin{pmatrix} ia & c+id \\ -c+id & -ia \end{pmatrix} \mid a,\ c,\ d \in \mathbf{R} \Big\},$$

so that a matrix-representation for the Adjoint representation of $\mathrm{SU}_2(\mathbf{C})$ on L_2 is a homomorphism

$$\mathrm{Ad}^{\boldsymbol{m}} : \mathrm{SU}_2(\mathbf{C}) \longrightarrow \mathrm{GL}_3(\mathbf{R}).$$

This "is" the desired projection, in the sense that it has kernel $\{\pm 1\}$, and image conjugate to $\mathrm{SO}_3(\mathbf{R})$ in $\mathrm{GL}_3(\mathbf{R})$ (depending on which basis of the Lie algebra L_2 is used to compute the matrix form of the representation).

In topological terms, the projection p is a non-trivial covering map of $\mathrm{SO}_3(\mathbf{R})$ (since $\mathrm{SU}_2(\mathbf{C})$ is connected). Thus $\mathrm{SO}_3(\mathbf{R})$ is *not simply connected* (in fact, one can show that $\mathrm{SU}_2(\mathbf{C})$ is simply connected, so it is the universal covering of $\mathrm{SO}_3(\mathbf{R})$). There are well-known "physical" demonstrations of this property of the rotation group (due in particular to Dirac); see, e.g., [**3**] for an accessible mathematical account, though seeing movies on the web might be even more enlightening...

6.4. The Hydrogen atom

We now come to the discussion of Example 1.2.3, i.e., of the basic invariants of simple quantum-mechanical systems, and in particular of the hydrogen atom.

In order to do this, we summarize briefly the fundamental formalism of (non-relativistic) quantum mechanics, constrasting it with classical newtonian mechanics, in the simplest situation of a single (point-like) particle evolving in $\mathbf{R}^3$, under the influence of some force (or forces):

- The *state* of the system at a given time t is represented by a unit vector ψ (i.e., with $\|\psi\| = 1$) in some fixed Hilbert space H [in contrast, in newtonian mechanics, the state of the particle is represented by an element $(x, p) \in \mathbf{R}^6$, where x represents the position of the particle and p its momentum $p = mv$, where $v \in \mathbf{R}^3$ is the the velocity at t and m is the mass of the particle];
- Two vectors ψ_1, ψ_2 in H correspond to the same state if and only if there exists $\theta \in \mathbf{R}$ such that $\psi_1 = e^{i\theta}\psi_2$, i.e., if the vectors are proportional;
- An *observable quantity* (or just "observable"), such as position or momentum, is represented by a linear operator A defined on a dense subspace D_A of H; if A is continuous, it can be defined on all of H, but many interesting observables are not continuous on D_A. Moreover A must be *self-adjoint*, which has the usual meaning when A is continuous on H, and has a more technical definition otherwise (see Exercise 6.4.1). [In Newtonian mechanics, an observable quantity is simply a real-valued function $f : P \longrightarrow \mathbf{R}$, where $P \subset \mathbf{R}^6$ is the set of possible states of the system.]
- The physical interaction of the system described by the state ψ with an observable A must result, though experiments, in some actual numerical (approximate) value; the crucial prediction of quantum mechanics is that this value λ will be an element of the *spectrum* $\sigma(A) \subset \mathbf{R}$ of A, but that it's value can *not* be predicted beforehand. Instead, one defines (purely mathematically) a probability measure $\mu_{\psi,A}$ on $\mathbf{R}$ such that $\mu_{\psi,A}(B)$ is the *probability* that the measurement will give a value in $B \subset \mathbf{R}$. The measure $\mu_{\psi,A}$ is called the *spectral measure of A with respect to v*. In the important case that A has (at most) countably many distinct eigenvalues $\lambda_i \in \mathbf{R}$, $i \geqslant 0$, whose eigenspaces span H, so that

$$H = \bigoplus_{i \geqslant 0} \ker(A - \lambda_i), \tag{6.16}$$

the spectral measure is defined by

$$\mu_{\psi,A}(B) = \sum_{\lambda_i \in B} \|p_i(\psi)\|^2, \tag{6.17}$$

where $p_i : H \longrightarrow \ker(A-\lambda_i)$ is the orthogonal projection on the i-th eigenspace. (This is a probability measure since ψ is a unit vector, and it satisfies the condition that any physically observed values would be among the eigenvalues λ_i of A, since the measure $\mu_{\psi,A}$ has support given by these eigenvalues.)

In particular, suppose A is a (non-trivial: $A \neq 0$, $A \neq \mathrm{Id}$) orthogonal projection. Its spectrum is $\{0, 1\}$, and the corresponding projections are just $p_1 = A$ itself and $p_0 = \mathrm{Id} - A$. Thus "measuring" the observable A will result in either of these values, with probability $\|A\psi\|^2$ of getting 1, and $1 - \|A\psi\|^2$ of getting 0 (in probabilistic terms, this is a Bernoulli random variable).

- The probability can be understood experimentally, and the prediction checked, as follows: if the measurement is repeated a large number of times (say N times), each one after preparing the system to be in state ψ, then the proportion N_B/N of the number N_B of measurements for which the experimental value λ is in B will be close to $\mu_{\psi,A}(B)$. [This is in striking contrast with newtonian mechanics: given that the particle is in the state $(x, p) \in P$, the value of the observation f is simply the exact value $f(x, p) \in \mathbf{R}$.] This property makes the link between the mathematical model and the natural world; it can, in principle, be falsified, but the probabilistic interpretation has turned out to be confirmed by test after test. Not only is it the case that the relative frequencies of various results (especially zero/one tests corresponding to projections) are found to be close to the theoretical values, but no method (either practical or even theoretical) has been found to predict exactly the values of the measurements one after the other.

  In the example where the spectral measure is given by (6.17), one will therefore "observe" the eigenvalue λ_i with relative frequency given by
  $$\mu_{\psi,A}(\{\lambda_i\}) = \|p_i(\psi)\|^2.$$
- Finally, the basic dynamical equation is Schrödinger's equation: there exists a particular observable E, the Hamiltonian, such that the state of the system evolves in time as a solution (ψ_t) of the equation
  $$i\frac{h}{2\pi}\frac{d}{dt}\psi_t = E\psi_t,$$
  here h is Planck's constant. The Hamiltonian encapsulates the forces acting on the particle. [In newtonian mechanics, the particle evolves according to the differential equation $m\frac{d^2}{dt^2}x =$ sum of the forces.] We won't discuss dynamics of quantum systems here, but it turns out that there is a connection between this equation and unitary representations of the (additive, non-compact) group $\mathbf{R}$; see Section 7.2.

For more information, written mostly from a mathematical point of view, the reader may refer to [**38, 41, 42, 44**] (there are also, of course, many physics books on quantum mechanics which may be worth reading.)

Exercise 6.4.1 (Unbounded self-adjoint operator). Let H be a Hilbert space, and let $A : D_A \longrightarrow H$ be a linear operator defined on $D_A \subset H$, a a dense subspace of H. The pair (D_A, A) is called an *unbounded* operator on H, and it is called *self-adjoint*, if the following two conditions hold: (1) we have
$$D_A = \{\psi \in H \mid \phi \mapsto \langle A\phi, \psi\rangle \text{ extends to a continuous linear form on } H\};$$
and (2) for all psi_1, $\psi_2 \in D_A$, we have
$$\langle A\psi_1, \psi_2\rangle = \langle\psi_1, A\psi_2\rangle.$$

Show that the following defines a self-adjoint unbounded operator:

$$\begin{aligned} H &= L^2(\mathbf{R}, dx) \\ D_A &= \{\psi \in H \mid x \mapsto x\psi(x) \in H\} \\ (A\psi)(x) &= x\psi(x) \text{ for } \psi \in D_A. \end{aligned}$$

This observable is interpreted as the *position* of a particle constrained to move on the line $\mathbf{R}$. Given $\psi \in L^2(\mathbf{R})$ with $\|\psi\|^2 = 1$, the measure $\mu_{\psi,A}$, in that case, is the probability measure $|\psi(x)|^2 dx$ on $\mathbf{R}$, so that the probability that a particle in the state described by ψ be located inside a set B is given by

$$\int_{\mathbf{R}} |\psi(x)|^2 dx.$$

Much more about the general theory of unbounded linear operators can be found, for instance, in the books of Reed and Simon [**31, 32**].

How does representation theory enter the picture? The answer has to do with possible symmetries of the system, which must be compatible with the linear structure underlying the Hilbert space involved. If an observable A of interest is also compatible with the symmetries of the system, and can be described using only eigenvalues (as in (6.16)), it follows that the eigenspaces must be invariant under these symmetries; in other words, if there is a symmetry group G of the system, the eigenspaces of observables are (unitary) *representations* of G.

Now consider such an eigenspace, say $V = \ker(A-\lambda)$. For states $\psi \in V$, the observable A has the specific, deterministic, value λ. If the representation V is *not* irreducible, we can find another observable B such that B commutes with A, and some eigenspace of B is a proper subspace W of V. For the states in W, both A and B have determined value. Thus, in the opposite direction, if V is an irreducible representation of G, nothing more may be said (deterministically) concerning the states in V.

What this shows is that, given a quantum system with symmetry group G, we should attempt to decompose the corresponding representation of G on H into irreducible representations. The states in each subrepresentation will be fundamental building blocks for all states (by linearity), which can not be further analyzed in a fully deterministic way.

We now illustrate these general principles with concrete examples. We consider a particle evolving in $\mathbf{R}^3$, which is constrained to lie on the unit sphere $\mathbf{S}^2 \subset \mathbf{R}^3$ (this restriction is not very physical, but it helps at first with the mathematical analysis, and there are many fascinating purely mathematical questions). The underlying Hilbert space is taken to be $H = L^2(\mathbf{S}^2, \nu)$, where ν is the surface Lebesgue measure on the sphere, defined similarly as in Example 5.2.4, (5). The operators of multiplication of a $\psi \in H$ by each coordinate function can play the role of position observables (as in Exercise 6.4.1), but since the coordinates are bounded, the corresponding operators are continuous on H). Suppose now that the system evolves according to a homogeneous force, compatible with the rotational symmetry of the sphere $\mathbf{S}^2$. The corresponding representation is therefore a representation of the rotation group $\mathrm{SO}_3(\mathbf{R})$, given by

$$(g \cdot \psi)(x) = \psi(g^{-1}x)$$

(this is indeed a unitary representation since the measure μ is $\mathrm{SO}_3(\mathbf{R})$-invariant; see Example 5.2.9).

The simplest observable to consider is the energy of the particle. In fact, even without knowing anything about its shape, it is intuitively clear that if the system is rotation-invariant, the energy must be compatible with the symmetries: in some sense, applying a rotation to a state ψ amounts to observing this state from a different direction in space, and rotation-invariance implies the absence of privileged directions!

Thus we attempt to decompose this representation of $\mathrm{SO}_3(\mathbf{R})$. Recall from Example 5.6.9 that the irreducible unitary representations of $\mathrm{SO}_3(\mathbf{R})$ are obtained using the projection

$$\mathrm{SU}_2(\mathbf{C}) \longrightarrow \mathrm{SO}_3(\mathbf{R})$$

from the odd-dimensional irreducible representations of $\mathrm{SU}_2(\mathbf{C})$: for each integer $\ell \geqslant 0$, there exists a unique irreducible representation of $\mathrm{SO}_3(\mathbf{R})$ of dimension $2\ell+1$, which we denote V_ℓ here.

Proposition 6.4.2 (Decomposition of $L^2(\mathbf{S}^2)$). *The space $L^2(\mathbf{S}^2)$ is isomorphic, as a representation of $\mathrm{SO}_3(\mathbf{R})$ to the Hilbert direct sum*

$$\bigoplus_{\ell \geqslant 0} V_\ell$$

of all irreducible representations of $\mathrm{SO}_3(\mathbf{R})$, each occuring with multiplicity 1.

Proof. There is a quick proof coming from Frobenius reciprocity, which starts from the observation that the group G acts transitively on $\mathbf{S}^2$, and the stabilizer of the point $\boldsymbol{n} = (1,0,0)$ is the subgroup $K \simeq \mathrm{SO}_2(\mathbf{R})$ of rotations around the x-coordinate axis.[5] Hence we have a bijection

$$\phi \begin{cases} K\backslash G & \longrightarrow & \mathbf{S}^2 \\ g & \mapsto & g \cdot \boldsymbol{n}, \end{cases}$$

which is in fact a homeomorphism (it is continuous, and both spaces are compact). It follows that functions on $\mathbf{S}^2$ are "the same" (by composition with this homeomorphism) as functions on G such that

$$f(kg) = f(g)$$

for $k \in K$, $g \in G$. In particular, the spaces $C(K\backslash G)$ and $C(\mathbf{S}^2)$ of continuous functions are isomorphic. Moreover, since the Lebesgue measure on $\mathbf{S}^2$ is known to be the unique rotation-invariant measure on $\mathbf{S}^2$, up to scalar, there exists a constant $c > 0$ such that

$$\phi_*\mu = c\nu.$$

We can therefore identify the representation on $L^2(\mathbf{S}^2, \nu)$ with the representation of $\mathrm{SO}_3(\mathbf{R})$ on

$$H_1 = \{f\,:\, \mathrm{SO}_3(\mathbf{R}) \longrightarrow \mathbf{C} \mid f(kg) = f(g) \text{ if } k \in K, g \in \mathrm{SO}_3(\mathbf{R})\}\}$$

with action

$$(g \cdot f)(x) = f(xg),$$

and the restriction of the usual inner product. This means that

$$L^2(\mathbf{S}^2, \nu) \simeq \mathrm{Ind}_K^{\mathrm{SO}_3(\mathbf{R})} \mathbf{1}$$

as unitary representation of $\mathrm{SO}_3(\mathbf{R})$ (see Example 5.2.10). Now, given an irreducible representation ϱ of $\mathrm{SO}_3(\mathbf{R})$, we can use the Frobenius Reciprocity formula for compact groups (Proposition 5.4.9) to derive

$$\dim \mathrm{Hom}_{\mathrm{SO}_3(\mathbf{R})}(\varrho, L^2(\mathbf{S}^2)) = \dim \mathrm{Hom}_K(\mathrm{Res}_K^{\mathrm{SO}_3(\mathbf{R})} \varrho, \mathbf{1}),$$

which is the multiplicity of the trivial representation in the restriction of ϱ to K. Now the point is that the diagonal subgroup in $\mathrm{SU}_2(\mathbf{C})$ maps onto K via the projection; although

[5] One could use any other point.

there are more intrinsic ways to see this, at least it can be checked using the "ugly" formula (5.24): for any $\theta \in \mathbf{R}$, we have

$$\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix} \mapsto \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos^2(\theta) - \sin^2(\theta) & -2\cos(\theta)\sin(\theta) \\ 0 & 2\cos(\theta)\sin(\theta) & \cos^2(\theta) - \sin^2(\theta) \end{pmatrix},$$

a rotation around the x-axis[6] with angle 2θ, and therefore this is also

$$\dim \mathrm{Hom}_T(\mathrm{Res}_T^{\mathrm{SU}_2(\mathbf{C})} \varrho, \mathbf{1}),$$

(seeing ϱ as a representation of $\mathrm{SU}_2(\mathbf{C})$). But we computed the restriction of the representations of $\mathrm{SU}_2(\mathbf{C})$ to the diagonal subgroup a long time ago: for $\varrho = \varrho_\ell$, we have

$$\mathrm{Res}_T^{\mathrm{SU}_2(\mathbf{C})} \varrho = \chi_{-2\ell} \oplus \chi_{-2\ell+2} \oplus \cdots \oplus \chi_{2\ell-2} \oplus \chi_{2\ell}$$

by (2.37) (remember that the m there is 2ℓ here) where

$$\chi_j\Bigl(\begin{pmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{pmatrix}\Bigr) = e^{ij\theta}.$$

By inspection, the multiplicity of the trivial representation $\mathbf{1} = \chi_0$ of T is indeed equal to 1... □

The physical meaning, for our hypothetical quantum particle on the sphere, is that if it is in a "pure" state ψ with well-defined energy, it has a natural invariant attached to it, namely the index ℓ such that ψ is in the subspace (say W_ℓ) of $L^2(\mathbf{S}^2)$ isomorphic to V_ℓ. This invariant is called the *azimuthal quantum number* of the state (or *orbital quantum number*).

Keeping with this particle on the sphere, suppose it is (i.e., the state ψ is) in W_ℓ. If we want to pinpoint the state more precisely, or at least describe specific states which can combine to construct all the states in W_ℓ, we must "break" the rotational symmetry.

Suppose we consider observables B which are only symmetric with respect to rotations around a fixed axis (say, the x-axis). This means that the underlying symmetry group becomes the subgroup $K \simeq \mathrm{SO}_2(\mathbf{R})$ of $\mathrm{SO}_3(\mathbf{R})$ of rotations around this axis. If we start from states known to be in W_ℓ, we must then decompose this space *as a representation of* K, and the corresponding K-subrepresentations represent states for which the energy and all K-invariant observables are fully known. Here K is abelian, so we know these subspaces are one-dimensional, and hence correspond to a unique state.

Precisely, as in the proof of Proposition 6.4.2, the space W_ℓ decomposes, as a representation of K, as the direct sum of the $2\ell + 1$ characters

$$\chi_j : \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \mapsto e^{ij\theta}$$

with $-\ell \leqslant j \leqslant \ell$ of $\mathrm{SO}_2(\mathbf{R})$. Each of the one-dimensional subspace $W_{\ell,j}$ on which K acts like χ_j therefore describes a unique quantum state, parameterized by the two quantum numbers $\ell \geqslant 0$ and $j \in \{-\ell, \ldots, \ell\}$. This second parameter is called the *magnetic quantum number* (historically, this is because it can be experimentally detected by putting systems in magnetic fields which are symmetric with respect to the given axis, for instance in the so-called "Zeeman effect".)

[6] This calculation depends on a compatible normalization of K and the projection, but changing either would just require to conjugate one or the other.

All this is still relevant for more realistic physical systems, where the particle evolves in $\mathbf{R}^3$, with state space given by $H = L^2(\mathbf{R}^3)$, under conditions of spherical symmetry, so that the relevant unitary representation of $\mathrm{SO}_3(\mathbf{R})$ is given by

$$\varrho(g)\varphi(x) = \varphi(g^{-1} \cdot x)$$

for $g \in \mathrm{SO}_3(\mathbf{R})$ and $\varphi \in L^2(\mathbf{R}^3)$. The point is that, as a representation of $\mathrm{SO}_3(\mathbf{R})$, we can separate the radius-dependency of functions in $L^2(\mathbf{R}^3)$ (on which $\mathrm{SO}_3(\mathbf{R})$ acts trivially) and the spherical components. To be precise:

Proposition 6.4.3. *There is a linear map*

$$C(\mathbf{S}^2) \otimes C_0([0, +\infty[) \xrightarrow{\Phi} L^2(\mathbf{R}^3)$$

mapping $\varphi \otimes \psi$ to the function

$$f(\boldsymbol{x}) = \psi(\|\boldsymbol{x}\|)\varphi\Bigl(\frac{\boldsymbol{x}}{\|\boldsymbol{x}\|}\Bigr).$$

This map is an isometry for the inner product induced by

$$\langle \varphi_1 \otimes \psi_1, \varphi_2 \otimes \psi_2 \rangle_0 = \Bigl(\int_0^{+\infty} \psi_1(r)\overline{\psi_2(r)}rdr\Bigr)\Bigl(\int_{\mathbf{S}^2} \varphi_1(x)\overline{\varphi_2(x)}d\nu(x)\Bigr)$$

and has dense image. Moreover

$$\varrho(g)f(x) = \Phi((g \cdot \varphi) \otimes \psi). \tag{6.18}$$

Sketch of proof. We leave again some details to the reader, but the main point is the spherical-coordinate integration formula, which states that

$$\int_{\mathbf{R}^3} f(x, y, z)dxdydz = \int_{\mathbf{S}^2} \int_0^{+\infty} f(rx)rdrd\nu(x)$$

for integrable $f : \mathbf{R}^3 \longrightarrow \mathbf{C}$. This implies in particular the isometry condition

$$\langle \Phi(\varphi_1 \otimes \Psi_1), \Phi(\varphi_2 \otimes \psi_2) \rangle_{L^2(\mathbf{R}^3)} = \langle \varphi_1 \otimes \Psi_1, \varphi_2 \otimes \psi_2 \rangle_0$$

for all $\varphi_i \in C(\mathbf{S}^2)$ and $\psi_i \in C_0([0, +\infty[)$. The formula (6.18) can be checked easily, and it remains to prove that the image of Φ is dense in $L^2(\mathbf{R}^3)$. But suppose f is orthogonal to this image, so that

$$\langle f, \Phi(\varphi \otimes \psi) \rangle = 0$$

for all φ and ψ. This translates, by the spherical-coordinates integration, to

$$0 = \int_{\mathbf{S}^2} \overline{\varphi(x)}\Bigl(\int_0^{+\infty} \overline{\psi(r)}f(rx)rdr\Bigr)d\nu(x)$$

for all $\varphi \in C(\mathbf{S}^2)$ and $\psi \in C_0([0, +\infty[)$. For all x, this gives

$$\int_0^{+\infty} \overline{\psi(r)}f(rx)rdr = 0$$

for all ψ, and then we get $f = 0$ (to be precise, one must invoke Fubini's Theorem since one really gets this for almost all $x \in \mathbf{S}^2$, and then $f(rx)$ for almost all $r \geqslant 0$, depending possibly on x...) □

In other words, we have

$$L^2(\mathbf{R}^3) \simeq L^2(\mathbf{S}^2, \nu)\hat{\otimes}L^2([0, +\infty[, rdr),$$

if we define the tensor product of Hilbert spaces on the right as the completion of $C(\mathbf{S}^2) \otimes C_0([0, +\infty[)$ with respect to the inner product $\langle \cdot, \cdot \rangle_0$, and this is an isomorphism as unitary $\mathrm{SO}_3(\mathbf{R})$-unitary representations, where the action of $L^2([0,+\infty, rdr)$ is trivial. Consequently, we obtain:

Corollary 6.4.4. *As a representation of* $\mathrm{SO}_3(\mathbf{R})$*, the space* $L^2(\mathbf{R}^3)$ *decomposes as a direct sum of infinitely many copies of every irreducible representation of* $\mathrm{SO}_3(\mathbf{R})$*.*

Thus, a state $f \in L^2(\mathbf{R}^3)$ lying in one irreducible subrepresentation still determines an azimuthal quantum number ℓ, and if the state is further compatible with breaking the radial symmetry as described above, it has a magnetic quantum number m, $-\ell \leqslant m \leqslant \ell$. But going further – pinpointing particular subspaces – requires more assumptions on the system, and is not purely a question of symmetry (i.e., of representation theory!) anymore...

CHAPTER 7

Other groups: a few examples

The picture of representation theory beyond the case of compact groups changes quite dramatically. Even if one restricts to locally compact groups, it happens that non-compact groups have typically infinite-dimensional irreducible unitary representations, and from these one can not usually produce all representations using Hilbert direct sums. Their study becomes in many respects more analytic, and certainly requires quite different tools and techniques. These are mostly beyond the scope of this book – as people say! – and the very modest goal of this chapter is simply to present some concrete examples of groups and, even if without full proofs, a survey of some phenomena involved in the representation theory of such groups, emphasizing mostly those that were not visible in the previous chapters.

We begin however with some words about a more algebraic topic concerning *algebraic groups.*

7.1. Algebraic groups

We have observed in the previous chapter that if $G \subset \mathrm{GL}_n(\mathbf{C})$ is a group of matrices which acts irreducibly on $\mathbf{C}^n$, one can not in general prove that

$$\dim B^G \leqslant \dim S^G$$

where B (resp. S) is the space of bilinear forms on $\mathbf{C}^n$ (resp. the space of hermitian forms). In particular, the inequality $\dim B^G \leqslant 1$, which holds for compact groups, can not be proved in the same manner using character theory. However, it is true in general; we will sketch here the proof of this fact, which can be seen a very simple introduction to *algebraic groups.*

Theorem 7.1.1. *Let G be an arbitrary group and let*

$$\varrho : G \longrightarrow \mathrm{GL}(V)$$

be a finite-dimensional irreducible complex representation of G. Then there is, up to scalar, at most one G-invariant bilinear form defined on V.

The reader could do worse at this point than trying to find a way to attack this question for $G = \mathrm{GL}_n(\mathbf{Z})$ for some $n \geqslant 2$ (matrices with integral coefficients and determinant ±1) with the "obvious" faithful representation of G on $V = \mathbf{C}^n$ (it is irreducible because, e.g., G contains permutation matrices, which act on $\mathbf{C}^n$ as the permutation representation of $\mathfrak{S}_n$; this representation of the symmetric group is not irreducible, but looking at Example 4.3.16 shows that a non-zero G-stable subspace $W \subset V$ which is not equal to W would be either the line spanned by $(1, \ldots, 1)$, or the space of vectors with sum of coordinates equal to zero, and it is a simple matter to check that neither of these subspace is stable under the whole of $\mathrm{GL}_n(\mathbf{Z})$.)

Proof. The first big idea of the proof is that one can "replace" the group G – about which nothing is assumed – by another one which is much nicer. As a first, easy, step, we can observe that a bilinear form b on V is G-invariant if and only if it is invariant under the image of ϱ in $\mathrm{GL}(V)$, and by replacing G by its image – which still acts irreducibly on V! –, we can assume that G is a subgroup of $\mathrm{GL}(V)$ (in other words, we can reduce to the case of a faithful irreducible representation.)

In particular, this means that the group – which may still be complicated – acts on V by matrix-vector multiplication. The crucial point for the next step is that this is a *polynomial* operation.

We are now going to replace G by a *larger* group $\mathbf{G} \subset \mathrm{GL}(V)$ – often, a much larger one – such that the G-invariant bilinear forms on G and on $\mathbf{G}$ coincide. This group $\mathbf{G}$ is known as the *Zariski-closure* of G in $\mathrm{GL}(V)$ (although Exercise 7.1.4 explains the terminology, which will also probably be familiar already to many readers, we present it here from scratch.)

Let B be the space of bilinear forms on V, on which we let $\mathrm{GL}(V)$ act in the usual way. Clearly, for any group $H \supset G$, the H-invariant bilinear forms $b \in B^H$ are also G-invariant. Hence the issue is to construct $\mathbf{G}$ so that any G-invariant bilinear form is also $\mathbf{G}$-invariant. Consider then a fixed G-invariant bilinear form b on V; if we further fix a basis (e_i) of V, the conditions

$$b(gv, gw) = b(v, w), \qquad \text{for all } g \in G,$$

can be summarized as the finitely many conditions $b(ge_i, ge_j) = b(e_i, e_j)$. We now look at those in terms of the coordinates $(g_{i,j})$ of $g \in \mathrm{GL}(V)$. Because G acts in the usual way on V, these amount to saying that, for all $g \in G$, the $(g_{k,\ell})$ satisfy certain *polynomial* relations, namely for all $g \in G$ and indices i, j, we have

$$\sum_{k,\ell}\sum b(e_k, e_\ell) g_{k,i} g_{\ell,j} - b(e_i, e_j) = 0 \tag{7.1}$$

(the various coefficients $b(e_k, e_\ell)$ and $b(e_i, e_j)$ are complex constants since b and the basis have been fixed.)

We now define two objects:

– the set of all polynomials vanishing on G, or more precisely, denoting

$$A = \mathbf{C}[(X_{i,j})_{i,j}, Y],$$

a polynomial algebra in $n^2 + 1$ variables,[1] we let

$$\mathcal{I}_G = \{f \in A \mid f(g_{i,j}, \det(g)^{-1}) = 0 \text{ for all } g = (g_{i,j}) \in G\} \quad ;$$

– the set of all elements in $\mathrm{GL}(V)$ which are common zeros of these polynomials, i.e

$$\mathcal{V}_G = \{x \in \mathrm{GL}(V) \mid f(x, \det(x)^{-1}) = 0 \text{ for all } f \in \mathcal{I}_G\}. \tag{7.2}$$

Obviously, the definitions show that $G \subset \mathcal{V}_G$. Moreover, since the invariance relations (7.1) – which are satisfied by G – are polynomial, they define elements $f_{i,j} \in \mathcal{I}_G$ (which happen to not depend on the extra variable Y). This means that those relations are also satisfied, by definition, by all elements $g \in \mathcal{V}_G$. In other words, the bilinear form b is invariant under all elements $g \in \mathcal{V}_G$.

The claim is now that this set $\mathcal{V}_G \subset \mathrm{GL}(V)$ is in fact a *subgroup* of $\mathrm{GL}(V)$. We will then denote $\mathbf{G} = \mathcal{V}_G$, and we have obtained the desired relation

$$B^G = B^{\mathbf{G}}. \tag{7.3}$$

[1] The usefulness of the presence of the extra variable Y, which is used to represent polynomially the inverse of a $g \in \mathrm{GL}(V)$, will be clear very soon.

So let us prove the claim, which is quite elementary. We start by showing that $\mathcal{V}_G$ is stable under inversion, since this is where the extra variable Y is useful. Given any $f \in \mathcal{I}_G$, define a new polynomial by

$$\tilde{f}((X_{i,j}), Y) = f(Y(\tilde{X}_{i,j}), \det(X_{i,j}))$$

where Y is seen as a scalar in the first argument, and the matrix $\tilde{X} = (\tilde{X}_{i,j})$ it multiplies is the comatrix of $X = (X_{i,j})$, i.e., the coefficients $\tilde{X}_{i,j}$ are the polynomials in $\mathbf{C}[(X_{i,j})]$ such that

$$\det(X) \times X^{-1} = \tilde{X}.$$

Thus $\tilde{f} \in A$; now, for $g \in \mathrm{GL}(V)$, we have

$$\tilde{f}(g, \det(g)^{-1}) = f(\det(g)^{-1}(\tilde{g}_{i,j}), \det(g)) = f(g^{-1}, \det(g)),$$

and this vanishes for all $g \in G$ since – because G is a group – $g^{-1} \in G$ and $f \in \mathcal{I}_G$. Thus $\tilde{f} \in \mathcal{I}_G$. This implies that $\tilde{f}$ vanishes on $\mathcal{V}_G$, i.e., for all $g \in \mathcal{V}_G$, we have

$$f(g^{-1}, \det(g)) = \tilde{f}(g, \det(g)^{-1}) = 0.$$

Finally, consider $g \in \mathcal{V}_G$ to be fixed; then $f(g^{-1}, \det(g)) = 0$ for all $f \in \mathcal{I}_G$, which by definition means that $g^{-1} \in \mathcal{V}_G$, as desired.

We now proceed to show that $\mathcal{V}_G$ is stable under products, along similar lines. First, we show that $\mathcal{V}_G$ is stable by multiplication by elements of G on both sides: if $g_1 \in G$ is given then for all $g \in \mathcal{V}_G$, both $g_1 g$ and $g g_1$ are in $\mathcal{V}_G$. Indeed, given $f \in \mathcal{I}_G$, we define

$$\tilde{f}(X_{i,j}, Y) = f(g_1 \cdot (X_{i,j}), \det(g_1)^{-1} Y)$$

where $g_1 \cdot (X_{i,j})$ denotes the matrix product. Since matrix multiplication is polynomial in the coordinates $(X_{i,j})$, this is an element of A. And as such, it belongs to $\mathcal{I}_G$, because for $g \in G$ we have $g_1 g \in G$ – since G is itself a group – and hence

$$\tilde{f}(g, \det(g)^{-1}) = f(g_1 g, \det(g_1 g)^{-1}) = 0$$

by definition of $\mathcal{I}_G$. Hence $\tilde{f}$ vanishes in fact identically on all of $\mathcal{V}_G$, which means that $f(g_1 g) = 0$ for all $g \in \mathcal{V}_g$. Since $f \in \mathcal{I}_F$ is arbitrary, this property applied to a fixed $g \in \mathcal{V}_G$ means that $g_1 g \in \mathcal{V}_G$. Reversing the order of a few products, we also obtain in this way that $g g_1 \in \mathcal{V}_G$ for all $g \in \mathcal{V}_G$.

Now we deduce the stability of $\mathcal{V}_G$ under all products: let $g_1 \in \mathcal{V}_G$ be given; for $f \in \mathcal{I}_G$, define again

$$\tilde{f}((X_{i,j}), Y)) = f(g_1 \cdot (X_{i,j}), \det(g_1)^{-1} Y) \in A.$$

We get $\tilde{f} \in \mathcal{I}_G$ (because, for $g \in G$, we know from above that $g_1 g \in \mathcal{V}_G$, and then $\tilde{f}(g, \det(g)^{-1}) = f(g_1 g, \det(g_1 g)^{-1}) = 0$, and therefore $\tilde{f}$ vanishes on $\mathcal{V}_G$; in particular, fixing some $g \in \mathcal{V}_G$ and using the fact that $f(g_1 g) = 0$ for all $f \in \mathcal{I}_G$, this means that $g_1 g \in \mathcal{V}_G$.

Now we continue the proof. Because of (7.3), we are left with having to prove

$$\dim B^{\mathbf{G}} \leqslant 1.$$

The reason this is easier, and why the passage from G to $\mathbf{G}$ is a drastic simplification (despite appearances at first sight!) is that $\mathbf{G}$ belongs to the category of *linear algebraic groups*, i.e., it is a subgroup of $\mathrm{GL}(V)$ which is the set of common zeros of a set of polynomial functions in A.

In fact, $\mathbf{G}$ is even better than a general algebraic group, because it is given with the inclusion $\mathbf{G} \subset \mathrm{GL}(V)$, which is a faithful irreducible (in particular, semisimple) representation (since $\mathbf{G} \supset G$, and G acts irreducibly on V, so does necessarily $\mathbf{G}$). An algebraic group of this type is is a *reductive group*.[2]

Now we must invoke a fairly deep fact without proof: if $\mathbf{G}$ is a reductive subgroup of $\mathrm{GL}(V)$, then it contains a *compact* subgroup $K \subset \mathbf{G}$ for which the Zariski-closure (computed by the same method as above, with $G = K$ instead) is still $\mathbf{G}$. Therefore we get

$$\dim B^G = \dim B^{\mathbf{G}} = \dim B^K \leqslant 1$$

by character theory applied to K as in the proof of Theorem 6.2.3 (see (6.4))... □

Example 7.1.2. One may ask why one does not try to go directly from G to the compact group K. This seems difficult because it may well happen that G and K have trivial intersection! A basic example is $G = \mathrm{GL}_n(\mathbf{Z})$, $n \geqslant 2$; then one can show (we will comment briefly on this in the next example) that $\mathbf{G}$ is the set of matrices g in $\mathrm{GL}_n(\mathbf{C})$ with $\det(g)^2 = 1$; the subgroup K is then the group of unitary matrices $g \in \mathrm{U}_n(\mathbf{C})$ with $\det(g)^2 = 1$. The intersection $G \cap K$ is the finite group W_n of signed permutation matrices (discussed in Exercise 4.7.13; indeed, if $g = (g_{i,j})$ is any unitary matrix with integral coefficients, the condition

$$\sum_j g_{i,j}^2 = 1$$

implies that, for a given i, a single $g_{i,j} \in \mathbf{Z}$ is ± 1, and the others are 0; denoting $j = \sigma(i)$, one sees that σ must be a permutation of $\{1, \ldots, n\}$, so that $g \in W_n$; since any $W_n \subset \mathrm{GL}_n(\mathbf{Z})$, we get the result.) In this case, $G \cap K$ still acts irreducibly on $\mathbf{C}^n$ (as the reader should check), but if we replace G by the subgroup

$$G_3 = \{g \in \mathrm{GL}_n(\mathbf{Z}) \mid g \equiv \mathrm{Id} \,(\mathrm{mod}\, 3)\},$$

which has finite index, it also possible to show that the Zariski-closure of G_3, and hence the compact subgroup K, are the same as that for G. However, $G_3 \cap K$ is now trivial since a signed permutation matrix which is congruent to the identity modulo 3 has to be the identity (we used reduction modulo 3 instead of 2 here to be able to distinguish the two signs.)

Example 7.1.3 (Examples of Zariski-closure). Going from G to the Zariski-closure $\mathbf{G}$ in the above proof might be a difficult psychological step at first, especially if one thinks of G as being particularly concrete (e.g., $G = \mathrm{GL}_n(\mathbf{Z}) \subset \mathrm{GL}_n(\mathbf{C})$) while $\mathbf{G}$ seems a very abstract object. However, it is a fact that computing the Zariski-closure of a group is quite often relatively easy, using known results (which may, of course, be quite deep and non-trivial.) Here are a few examples.

Exercise 7.1.4 (The Zariski topology). The association of the "big" group $\mathbf{G}$ to the group G has the aspect of a "closure" operation. Indeed, it can be interpreted as taking the closure of G in $\mathrm{GL}(V)$ with respect to a certain topology, called the Zariski topology.

Let k be an algebraically closed field and let $n \geqslant 1$ be an integer. We define $\mathbf{A}^n = k^n$, which is called the affine n-space over k, and $A = k[X_1, \ldots, X_n]$. Note that for a polynomial $f \in A$ and $x \in \mathbf{A}^n$, we can evaluate f at x.

(1) For any subset $\mathcal{I} \subset A$, let

$$\mathcal{V}(\mathcal{I}) = \{x \in \mathbf{A}^n \mid f(x) = 0 \text{ for all } f \in \mathcal{I}\}.$$

[2] One definition of a general reductive group is that it is an algebraic group which has a finite-dimensional semisimple representation ϱ with *finite* kernel.

Show that the collection of sets $(\mathcal{V}(\mathcal{I}))_{\mathcal{I}\subset A}$ is the collection of *closed* sets for a topology on $\mathbf{A}^n$.

In particular, this allows us to speak of the *Zariski-closure* of any subset $V \subset \mathbf{A}^n$: it is the intersection of all Zariski-closed subsets of $\mathbf{A}^n$ which contain V.

(2) For $n = 1$, show that a subset $V \subset \mathbf{A}^1$ is closed for the Zariski topology if and only if $V = \mathbf{A}^1$ or V is finite. Show that the Zariski topology on $\mathbf{A}^2$ is *not* the product topology on $\mathbf{A}^1 \times \mathbf{A}^1$. Show also that the Zariski topology is *not* Hausdorff, at least for $\mathbf{A}^1$ (the case of $\mathbf{A}^n$ for arbitrary n might be more difficult.)

(3) Let $G = \mathrm{GL}_n(k)$, seen as a subset of $\mathbf{A}^{n^2}$ by means of the matrix coefficients. Show that G is dense in $\mathbf{A}^{n^2}$ with respect to the Zariski topology. Furthermore, show that the set

$$\tilde{G} = \{(g, x) \in \mathbf{A}^{n^2+1} \mid \det(g)x = 1\} \subset \mathbf{A}^{n^2+1}$$

is Zariski-closed in $\mathbf{A}^{n^2+1}$.

(4) For $k = \mathbf{C}$ and $G \subset \mathrm{GL}_n(\mathbf{C})$, show that the Zariski-closure of $\tilde{G} \subset \mathbf{A}^{n^2+1}$ is equal to

$$\tilde{\mathbf{G}} = \{(g, x) \in \mathbf{G} \times \mathbf{C} \mid \det(g)x = 1\}.$$

The Zariski topology is a foundational notion in algebraic geometry; readers interested in learning more may look at [**18**] for the general theory (or should really try to attend a course in algebraic geometry, if at allpossible!). In the context of linear algebraic groups, the books [**4**] and [**40**] mayb be more accessible.

Exercise 7.1.5 (Zariski closure and polynomial representations). Let k be an algebraically closed field and let $G \subset \mathrm{GL}_n(k)$ be any subgroup. Let

$$\varrho : G \longrightarrow \mathrm{GL}_m(k)$$

be a (matrix) representation of G. We assume that ϱ is *polynomial*, i.e., that the matrix coefficients of ϱ are functions on G which are restrictions of polynomials in the coordinates $g_{i,j}$ and in $\det(g)^{-1}$.

(1) Define $\mathbf{G} \subset \mathrm{GL}_n(k)$ as the Zariski closure of G. Show that there exists a unique representation

$$\varrho_1 : \mathbf{G} \longrightarrow \mathrm{GL}_m(k)$$

such that ϱ coincides with ϱ on G. What is ϱ_1 if ϱ is the injection of G in $\mathrm{GL}_n(k)$?

(2) Show that a subspace $V \subset k^m$ is a subrepresentation of ϱ if and only if it is a subrepresentation of ϱ_1.

(3) Show that the subspace $(k^m)^G$ of vectors invariant under G is equal to the subspace of vectors invariant under $\mathbf{G}$.

(4) Show that the representations $\tilde{\varrho}$ (contragredient), $\mathrm{Sym}^k \varrho$ (k-th symmetric power, for $k \geqslant 0$), $\bigwedge^k \varrho$ (alternating power, for $k \geqslant 0$) are also polynomial.

Remark 7.1.6 (The Larsen alternative). The Larsen Alternative can also be phrased in terms of algebraic groups, if the fourth moment invariant $\mathrm{M}_4(G)$ of a subgroup $G \subset \mathrm{GL}_n(\mathbf{C})$ is defined algebraically as the dimension of the space of G-invariants in $\mathrm{End}(\mathbf{C}^n) \otimes \mathrm{End}(\mathbf{C}^n)$, as in Lemma 6.3.3. The statement is the following (see [**22**, Th. 1.1.6 (1)]):

Theorem 7.1.7 (Larsen). *Let $G \subset \mathrm{GL}_n(\mathbf{C})$ be a reductive linear algebraic group. If $\mathrm{M}_4(G) = 2$, then either G is finite, or $\mathrm{SL}_n(\mathbf{C}) \subset G$.*

We now sketch another application of algebraic groups to a basic question of representation theory, which is due to Chevalley: over $\mathbf{C}$, the tensor product of two semisimple representations remains semisimple (with no continuity or related assumption!)

THEOREM 7.1.8 (Chevalley). *Let ϱ_1, ϱ_2 be finite-dimensional complex semisimple representations of a group G. Then $\varrho_1 \otimes \varrho_2$ is semisimple.*

As mentioned by Serre [**36**, Part II, §1], it would be interesting to have an "elementary" proof of this fact, not involving algebraic groups.

SKETCH OF PROOF. □

7.2. Locally-compact abelian groups

7.3. A non-abelian example: $\mathrm{SL}_2(\mathbf{R})$

APPENDIX A

Some useful facts

A.1. Algebraic integers

Readers familiar with algebraic integers will probably not need any information from this section. We recall the definition and some properties of algebraic integers which are relevant to the applications in the text, especially in Section 4.7.2. A very good summary for similar purposes (with proofs) is found in Section 4.3 of [**11**].

An algebraic integer $z \in \mathbf{C}$ is any complex number such that there exists a non-zero *monic* polynomial $p \in \mathbf{Z}[X]$ such that $p(z) = 0$; we denote by $\bar{\mathbf{Z}}$ the set of algebraic integers.[1] Since the set $I_z = \{p \in \mathbf{Z}[X] \,:\, p(z) = 0\}$ is an ideal of $\mathbf{Z}[X]$, there exists a monic generator p_z, which is called the *minimal polynomial* of z.

For instance, any $n \in \mathbf{Z}$ is a zero of $p = X - n$, and hence $\mathbf{Z} \subset \bar{\mathbf{Z}}$. In fact, we have the following stronger fact, which illustrates one way in which algebraic integers generalize integers:

Proposition A.1.1. *An algebraic integer z which is also a rational number is in fact an element of $\mathbf{Z}$, i.e., we have $\bar{\mathbf{Z}} \cap \mathbf{Q} = \mathbf{Z}$.*

Proof. To see this, let $z = a/b$, with a and b coprime integers, be an algebraic integer, so that we have

$$X^n + a_{n-1}X^{n-1} + \cdots + a_1X + a_0 = 0,$$

for some $n \geqslant 1$ and integral coefficients $a_i \in \mathbf{Z}$.

Substituting the value of z and multiplying through with the common denominator b^n, one finds

$$a^n + a_{n-1}a^{n-1}b + \cdots + a_1ab^{n-1} + a_0b^n = 0,$$

and therefore $b \mid a^n$, which means $b = 1$ or -1 since $(a, b) = 1$. □

Other important examples of algebraic integers include arbitrary roots of unity (solutions of $X^n - 1 = 0$) Moreover, although this is not entirely obvious, $\bar{\mathbf{Z}}$ is a ring: the sum, different and product of algebraic integers remains an integer. We prove this in an ad-hoc, yet fairly elegant manner (this is the approach used in [**11**]):

Proposition A.1.2. (1) *A complex number $z \in \mathbf{C}$ is an algebraic integer if and only there exists a square matrix $A \in \mathrm{M}_n(\mathbf{Z})$, for some $n \geqslant 1$, with integral coefficients, such that z is an eigenvalue of A, i.e., such that* $\det(z - A) = 0$.

2) *If z_1, z_2 are algebraic integers, then so are $z_1 + z_2$ and z_1z_2.*

Proof. (1) The characteristic polynomial $\det(X - A)$ of A is a monic integral polynomial of degree n, and therefore any of its zeros is an algebraic integer. We must see the converse. Thus let z be an algebraic integer, and let

$$p = X^n + a_{n-1}X^{n-1} + \cdots + a_1X + a_0 in\mathbf{Z}[X]$$

[1] It is the restriction to monic equations which is crucial to obtain a generalization of integers; if instead p is allowed to be any non-zero $p \in \mathbf{Q}[X]$, one obtains the notion of *algebraic number*.

be the minimal polynomial of z. We can write down immediately a suitable matrix, namely

$$A_z = \begin{pmatrix} 0 & \cdots & \cdots & \cdots & -a_0 \\ 1 & 0 & \cdots & \cdots & -a_1 \\ 0 & 1 & 0 & \cdots & -a_2 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & 1 & -a_{n-1} \end{pmatrix}.$$

It is a standard exercise that $\det(A_z - z) = 0$, but the following explains how to "produce" it: consider the abelian group $M = \mathbf{Z}[z]$ generated by powers of z, as a subgroup of $\mathbf{C}$. Because of the relation $p(z) = 0$, this is a free finitely-generated abelian group, with basis

$$(1, z, z^2, \ldots, z^{n-1})$$

hence of rank n. The map

$$m_z \begin{cases} M & \longrightarrow & M \\ a & \mapsto & za \end{cases}$$

is a homomorphism of abelian groups, and with respect to the basis of M above, it is represented by an integral matrix, which is precisely A_z. But since $m_z(z) = z^2 = z \times z$, we see that z is an eigenvalue of the homomorphism m_z, hence of the matrix A.

(2) We can now prove that $\bar{\mathbf{Z}}$ is stable under product and multiplication using the characterization we just obtained. Indeed, let A_1 and A_2 be integral matrices such that z_i is an eigenvalue of A_i. It is standard that z_1z_2 is an eigenvalue of $A_1 \otimes A_2$, which is also an integral matrix (if e_i are eigenvectors of A_i for the eigenvalue z_i, then $(A_1 \otimes A_2)(e_1 \otimes e_2) = z_1z_2(e_1 \otimes e_2)$, by definition of tensor products.)

The case of sums is a bit less obvious, and the formula is worth remembering: $z_1 + z_2$ is an eigenvalue of $A_1 \otimes \mathrm{Id} + \mathrm{Id} \otimes A_2$. Indeed, with e_i as before, we have

$$(A_1 \otimes \mathrm{Id} + \mathrm{Id} \otimes A_2)(e_1 \otimes e_2) = z_1(e_1 \otimes e_2) + z_2(e_1 \otimes e_2),$$

so $e_1 \otimes e_2$ is an eigenvector for the eigenvalue $z_1 + z_2$. □

We can now see that any linear combination of roots of unity with *integral* coefficients is an algebraic integer. In particular:

Corollary A.1.3. *Let G be a finite group, and let ϱ be a finite-dimensional complex representation of G. Then for any $g \in G$, the character value $\chi_\varrho(g)$ is an algebraic integer.*

Proof. Indeed, $\chi_\varrho(g) = \operatorname{Tr} \varrho(g)$ is the sum of the $\dim(\varrho)$ eigenvalues of $\varrho(g)$, and each of them is a root of unity, since $\varrho(g)^{|G|} = 1$ (using the finiteness of G). □

We now discuss quickly the divisibility relation in $\bar{\mathbf{Z}}$. As might be expected, if z_1, z_2 are algebraic integers, one says that z_1 divides z_2, denoted

$$z_1 \mid z_2,$$

if and only if $z_2 = z_1z$ with z also an algebraic integer (in other words, the ratio $z_2/z_1 \in \mathbf{C}$ is in fact in $\bar{\mathbf{Z}}$).

We see clearly that if the same z_1 divides z_2 and z_3 (all in $\bar{\mathbf{Z}}$), it divides their sum or difference, or their product. In particular, if we have a relation

$$\sum_i z_iw_i = 1$$

with z_i, w_i all algebraic integers, and fix some natural integer $q \geqslant 2$, we can conclude that some w_i is *not* divisible by q, in view of the fact that $1/q \notin \bar{\mathbf{Z}}$.

One can also define a coprimality relation between algebraic integers: if z_1 and z_2 are algebraic integers, they are said to be coprime, which is denoted $(z_1, z_2) = 1$, if and only if there exist algebraic integers w_1, w_2 such that

$$z_1 w_1 + z_2 w_2 = 1.$$

This shows in particular that if two ordinary integers in $\mathbf{Z}$ are coprime (in the usual sense), they are also coprime as algebraic integers. On the other hand, suppose z_1 and z_2 have "common divisor" $w \in \bar{\mathbf{Z}}$, i.e., we have $w \mid z_1$ and $w \mid z_2$. Then the algebraic integers z_1 and z_2 can be coprime only if w is a *unit*, i.e., if $1/w$ is also in $\bar{\mathbf{Z}}$. Indeed, writing $z_i = wy_i$, we get

$$wy_1 w_1 + wy_2 w_2 = 1$$

and thus $1/w = y_1 w_1 + y_2 w_2 \in \bar{\mathbf{Z}}$. If $w \in \mathbf{Z}$, of course, the condition $1/w \in \bar{\mathbf{Z}}$ means that $w = \pm 1$.

Remark A.1.4 (Units). There are many examples of units, some of which are complex numbers with modulus not equal to 1, in contrast with the units ±1 in $\mathbf{Z}$. For instance, the element $\varepsilon = 1 + \sqrt{2}$ satisfies

$$\frac{1}{1+\sqrt{2}} = \frac{1-\sqrt{2}}{1-2} = -1 + \sqrt{2} \in \bar{\mathbf{Z}},$$

so it is a unit, although $|\varepsilon| > 1$.

Other examples of units are roots of unity, since the inverse of a root of unity is also one.

The fundamental link between divisibility and coprimality remains valid:

Proposition A.1.5. *Let z_1, z_2, z_3 be algebaic integers such that $z_1 \mid z_2 z_3$. If z_1 and z_2 are coprime, then z_1 divides z_3.*

Proof. Indeed, from a relation

$$z_1 w_1 + z_2 w_2 = 1$$

we get $z_3 = z_1 w_1 z_3 + z_2 z_3 w_2$, and since each term of the sum is divisible by z_1, so is z_3. □

The last useful topic is the definition of algebraic conjugates of an algebraic integer, and of the norm map. The former is very natural:

Definition A.1.6 (Conjugates and norm of algebraic integers). Let $z \in \bar{\mathbf{Z}}$ be an algebraic integer. Let $p \in \mathbf{Z}[X]$ be the *unique* monic polynomial with integral coefficients which is irreducible over $\mathbf{Q}$ and such that $p(z) = 0$. A *conjugate* w of z is any root of p in $\mathbf{C}$, and the *norm* of z, denoted $N(z)$, is the product of all conjugates of z.

In other words, a conjugate of z is an algebraic integer which satisfies "the same equation" as z, where one normalizes the monic polynomial of which z is a root by taking the one of smallest possible degree. Since any polynomial $p_1 \in \mathbf{Z}[X]$ for which $p_1(z) = 0$ is a multiple of p, it follows that whenever z satisfies a polynomial relation with integral coefficients, so do its conjugates.

For example: if $n \in \mathbf{Z}$, we have $p = X - n$, and n has no other conjugate than itself; if z is a root of unity, then all its conjugates are also roots of unity (but not all roots of unity are conjugates of z, e.g., -1 is not a conjugate of $z = i$, because the minimal polynomial for i is $X^2 + 1$, and not $X^4 - 1$.)

Factoring the minimal polynomial p, we have

$$p(X) = \prod_w (X - w)$$

where the product, by definition, is over all conjugates of z. In particular, we find

$$N(z) = \prod_w w = (-1)^n p(0).$$

with $n = \deg(p)$. In particular, we see that the norm of z is an integer in $\mathbf{Z}$.

Proposition A.1.7 (Conjugates of sums and products). *Let z_1 and z_2 be algebraic integers. Then any conjugate w of $z_1 z_2$ can be written*

$$w = w_1 w_2$$

with w_1 a conjugate of z_1 and w_2 a conjugate of z_2. Similarly, any conjugate of $z_1 + z_2$ is of the form $w_1 + w_2$ for some conjugate w_i of z_i.

Note, however, that not *all* sums $w_1 + w_2$ are necessarily conjugates of $z_1 + z_2$ (for instance, take $z_1 = \sqrt{2}$, $z_2 = -\sqrt{2}$; then $w_1 = z_1$ and $w_2 = \sqrt{2}$ are conjugates of z_1 and z_2, with $w_1 + w_2 = 2\sqrt{2}$, although $z_1 + z_2 = 0$ has no non-zero conjugate.)

Proof. Although this property is much better understood in terms of Galois theory, there is a cute argument using the criterion in Proposition A.1.2. We present this for $z_1 z_2$, leaving the case of the sum to the reader.

Consider integral matrices A_1 and A_2 with characteristic polynomials p_1 and p_2, the minimal polynomials for z_1 and z_2 respectively. Form the matrix $A = A_1 \otimes A_2$. Since $z_1 z_2$ is an eigenvalue of A, we know that any conjugate of $z_1 z_2$ is among the other eigenvalues of A. Similarly, for any conjugate w_1 of z_1 and w_2 of z_2, there are eigenvectors e_1, e_2 with $A_i e_i = w_i e_i$. Each of these gives an eigenvector $e_1 \otimes e_2$ of A with eigenvalue $w_1 w_2$. If we count, we see that we construct this way $n_1 n_2$ eigenvectors of A, with eigenvalues given by products of conjugates of z_1 and z_2. Since A has size $n_1 n_2$, there can be no other eigenvector, and therefore no other eigenvalue either! □

Corollary A.1.8. *Let z be an algebraic integer and $n \in \mathbf{Z}$ such that $n \mid z$. Then n^r divides $N(z)$ in the ring $\mathbf{Z}$, where r is the degree of the minimal polymomial of z, or in other words the number of conjugates of z.*

Proof. The point is that n divides any conjugate w of z: indeed, if we write $z = nz_1$ for some algebraic integer z_1, we see that w, as a conjugate of nz_1, must be of the form nw_1, where w_1 is a conjugate of z_1 (since n is the only conjugate of itself...) □

Exercise A.1.9. Let z be an algebraic integer. Show that z is a unit in $\bar{\mathbf{Z}}$ if and only if $N(z) = \pm 1$.

A.2. The spectral theorem

In the proof of the general case of the Peter-Weyl theorem, a crucial ingredient is the fact that certain operators constructed using the regular (or left-regular) representation have non-trivial, but finite-dimensional, eigenspaces. The standard statement along these lines is the spectral theorem for compact normal operators. We will state this result, but we will only prove a weaker statement that is sufficient for our purposes.

We start with the definition:

Definition A.2.1 (Compact operator). Let H be a Hilbert space and let $T : H \longrightarrow H$ be a continuous linear operator. Then T is *compact* if and only if there is a sequence (T_n) of continuous operators $T_n : H \to H$ such that $\dim \mathrm{Im}(T_n) < +\infty$ for all n, and $T_n \to T$ in the operator norm topology, i.e., uniformly on the unit ball of H.

Example A.2.2. A compact operator should be considered as "small" in some sense. An illustration of this intuitive idea is the fact that if $\lambda \neq 0$, the operator $\lambda\mathrm{Id}$ on H is compact if and only if H is finite-dimensional. Indeed, suppose $T_n \to \lambda\mathrm{Id}$ in $\mathrm{L}(H)$. Then, for n sufficiently large, we $\|\lambda^{-1}T_n - \mathrm{Id}\| < 1$, and this implies that $\lambda^{-1}T_n$ is invertible (this is well-known: check that the geometric series

$$\sum_{k\geqslant 0} (\mathrm{Id} - \lambda^{-1}T_n)^k$$

converges in $\mathrm{L}(H)$, as it should, to the inverse of $\mathrm{Id} - (\mathrm{Id} - \lambda^{-1}T_n) = \lambda^{-1}T_n$). Then

$$\dim H = \dim \mathrm{Im}(T_n) < +\infty.$$

The spectral theorem is the following:

Theorem A.2.3 (Spectral theorem). *Let H be a Hilbert space and let $T : H \longrightarrow H$ be a normal compact operator, i.e., such that $TT^* = T^*T$, for instance a self-adjoint operator with $T = T^*$. Then there exists a subset $S \subset \mathbf{C}$ which is finite or countable, such that the eigenspace* $\ker(T - \lambda)$ *is non-zero if and only if $\lambda \in S$, and is finite-dimensional if $\lambda \neq 0$, and such that furthermore we have a Hilbert space orthogonal direct sum decomposition*

$$H = \bigoplus_{\lambda \in S} \ker(T - \lambda).$$

In particular, if $T \neq 0$, the set S is not reduced to $\{0\}$.

In other words: a normal compact operator can be diagonalized with at most countable countably many eigenvalues, and for any non-zero eigenvalue, the corresponding eigenspace is finite-dimensional. Note that, if T has finite rank (i.e., $\dim \mathrm{Im}(T) < +\infty$), this is just a form of the spectral theorem for matrices of finite size commuting with their adjoint. In view of the definition of compact operators, the result can therefore be seen as "passing to the limit" with this classical theorem.

We will apply this to Hilbert-Schmidt integral operators:

Proposition A.2.4. *Let $(X, \mu)X$ be a measure space*

$$k : X \times X \longrightarrow \mathbf{C}$$

be a function which is in $L^2(X \times X, \mu \times \mu)$. The linear operator T_k acting on $L^2(X, \mu)$ by

$$(T_k\varphi)(x) = \int_X k(x, y)f(y)d\mu(y)$$

is compact. It is self-adjoint if k satisfies $k(x, y) = \overline{k(y, x)}$.

The operator T_k is customarily called the *Hilbert-Schmidt (integral) operator with kernel k*. We give the proof in the case when $L^2(X, \mu)$ is a separable Hilbert space, which is the case in most applications (for instance for X a compact metric space and μ a Radon measure on X). The general case is considered, e.g., in [**10**, XI.8.44].

Proof when $L^2(X,\mu)$ is separable. First of all, the Cauchy-Schwarz inequality and Fubini's theorem give

$$\int_X |T_k(\varphi)(x)|^2 d\mu(x) \leqslant \int_X \Bigl(\int_X |k(x,y)|^2 d\mu(y)\Bigr)\Bigl(\int_X |f(y)|^2 d\mu(y)\Bigr) d\mu(x)$$
$$= \|k\|^2_{L^2(X\times X)} \|f\|^2_{L^2(X)},$$

which shows that T_k is well-defined, and continuous with norm

$$\|T_k\| \leqslant \|k\|_{L^2(X\times X)}. \tag{A.1}$$

We can now check quite easily that T_k is compact when $L^2(X,\mu)$ is separable. Fix any orthonormal basis $(\varphi_k)_{k\geqslant 1}$ of $L^2(X)$. Then the functions

$$\psi_{k,\ell} : (x,y) \mapsto \varphi_k(x)\varphi_\ell(y)$$

are known to form an orthonormal basis of $L^2(X\times X)$ (see, e.g., [**31**, II.4, p. 51]), and we can therefore expand k in an L^2-convergent series

$$k = \sum_k \sum_\ell \alpha(k,\ell)\psi_{k,\ell}, \qquad \alpha(k,\ell) \in \mathbf{C}. \tag{A.2}$$

Given any $N \geqslant 1$, we define the approximation $T^{(N)} = T_{k_N}$, where k_N is the corresponding partial sum

$$k_N = \sum_{k\leqslant N} \sum_{\ell\leqslant N} \alpha(k,\ell)\psi_{k,\ell}.$$

Note that for any $\varphi \in L^2(X)$, we have

$$T^{(N)}\varphi = \sum_{k\leqslant N} \sum_{\ell\leqslant N} \alpha(k,\ell)\Bigl(\int_X \varphi_\ell(y) f(y) d\mu(y)\Bigr)\varphi_k,$$

which shows that the image of $T^{(N)}$ is finite-dimensional, spanned by the φ_k, $k \leqslant N$. In addition, in the space of operators on $L^2(X)$, we have

$$\|T_k - T^{(N)}\| = \|T_{k-k_N}\| \leqslant \|k - k_N\|_{L^2(X\times X)},$$

by (A.1), which tends to 0 as $N \to +\infty$ (since the series (A.2) converges in $L^2(X\times X)$), we have found finite rank approximations of T_k, and thus T_k is compact.

Finally, it is a formal computation, left to the reader, to check that

$$T_k^* = T_{\tilde{k}}, \qquad \tilde{k}(x,y) = \overline{k(y,x)},$$

so that k is self-adjoint if $k = \tilde{k}$. □

We are now going to prove part of the spectral theorem in a special case involving Hilbert-Schmidt operators. This statement is enough for the application to the proof of the Peter-Weyl theorem.

Proposition A.2.5. *Let X be a compact space with a Radon measure μ, and let $T = T_k$ be a non-zero self-adjoint Hilbert-Schmidt operator with continuous kernel k such that T_k is non-negative, i.e.*

$$\langle T_k f, f\rangle \geqslant 0$$

for all $f \in L^2(X,\mu)$.

Then there exists a positive eigenvalue $\lambda > 0$ of T, and the corresponding eigenspace is finite-dimensional.

PROOF WHEN $L^2(X,\mu)$ IS SEPARABLE. We denote $H = L^2(X)$ in this argument, since parts of it will apply to any positive compact operator. We first show the general fact that if $\lambda \neq 0$ is an eigenvalue of a compact operator T, the eigenspace $V = \ker(T-\lambda)$ is finite-dimensional (this is false for $\lambda = 0$, e.g., take $T = 0$...). The basic idea is that T, restricted to V, remains compact – but this operator on V is $\lambda\mathrm{Id}$, and we can apply Example A.2.2. To implement this, let P denote the orthogonal projector onto the closed subspace V. Consider a sequence (T_n) of operators such that $T_n \to T$ in $\mathrm{L}(H)$ and $\mathrm{Im}(T_n)$ is finite-dimensional, and define $U_n = PT_n$ as operator on V. These are finite rank operators in $\mathrm{L}(V)$, and we have

$$\|U_n - T\|_{\mathrm{L}(V)} \leqslant \|(P-\mathrm{Id})T_n\|_{\mathrm{L}(V)} + \|T_n - T\|_{\mathrm{L}(H)} = \|T_n - T\|_{\mathrm{L}(H)}$$

since $P = \mathrm{Id}$ on V. This shows that, indeed, T is compact when restricted to the stable subspace V.

Now for the existence of the eigenvalue, which is the most crucial part. The idea is that the norm $\lambda = \|T\|_{\mathrm{L}(H)}$ itself is an eigenvalue, and in order to prove this it is enough to show that the supremum

$$\|T\|_{\mathrm{L}(H)} = \sup_{v\neq 0} \frac{\|Tv\|}{\|v\|}$$

is reached. Indeed, if $v \in H$ has norm 1 and satisfies $\|Tv\| = \lambda$, consider the functions

$$\alpha_w \,:\, t \mapsto \frac{\|T(v+tw)\|^2}{\|v+tw\|^2}$$

for fixed $w \in H$ and $t \in \mathbf{R}$ close to 0. By definition of the supremum, the point $t = 0$ is an extremum, and hence $\alpha'_w(0) = 0$. But we can compute this: writing

$$\alpha_w(t) = \frac{\|Tv\|^2 + t^2\|Tw\|^2 + 2t\,\mathrm{Re}\langle Tv, Tw\rangle}{\|v\|^2 + t^2\|w\|^2 + 2t\,\mathrm{Re}\langle v, w\rangle} = \frac{\lambda^2 + at^2 + 2bt}{1 + ct^2 + 2dt}$$

(say, which incidentally confirms that α_w is differentiable), we obtain

$$\alpha'_w(t) = \frac{(2at+2b)(1+ct^2+2dt) - (\lambda^2+at^2+2bt)(2ct+2d)}{(1+ct^2+2dt)^2},$$

and hence[2]

$$\alpha'_w(0) = 2b - 2d\lambda^2 = 2\,\mathrm{Re}\langle Tv, Tw\rangle - 2\lambda^2\,\mathrm{Re}\langle v, w\rangle.$$

In other words, since T is self-adjoint and λ is real, we have

$$\mathrm{Re}\langle T^2 v, w\rangle = \mathrm{Re}\langle Tv, Tw\rangle = \mathrm{Re}\langle \lambda^2 v, w\rangle$$

for all $w \in H$. Taking $w = T^2v - \lambda^2 v$, we obtain $T^2 v = \lambda^2 v$. Now we "take the squareroot" as follows: we write $Tv = \lambda v + v_1$, and apply T, getting

$$\lambda^2 v = T^2 v = \lambda T v + T v_1 = \lambda^2 v + \lambda v_1 + T v_1,$$

so $Tv_1 = -\lambda v_1$. But $\lambda > 0$ and T is positive, so this is impossible unless $v_1 = 0$, and hence $Tv = \lambda v$, as desired.

Thus we are reduced to proving the existence of a vector achieving the norm of T (indeed, this reduction did not use anything about T itself!) Here the idea is that this property holds for a finite-rank operator – and we hope to get it to carry through the limiting process! Readers familiar with weak convergence and the alternative definition of compact operators based on the relative compactness of the image of the unit ball will

[2] This can be obtained also by looking at the inequality $\|T(v+tw)\|^2 \leqslant \lambda^2\|v+tw\|^2$

find the following arguments rather naïve, but one should remember that the Peter-Weyl theorem predates such notions.

Thus we argue from scratch, starting with a sequence (v_n) of unit eigenvectors of T_n for the norm $\lambda_n = \|T_n\|$. Since $\lambda_n \to \lambda$, it would of course be enough to know that (v_n), or a subsequence of (v_n), converges. But since the unit ball of an infinite-dimensional Hilbert space is not compact, we can not claim that such a subsequence exists. However, we can fix a countable dense subset (φ_k) of H, consisting of the $\mathbf{Q}$-linear span of an orthonormal basis (ψ_m) of H. Then, by a diagonal argument, we can find a subsequence $w_j = v_{n_j}$ of (v_n) such that

$$\langle \varphi_k, w_j \rangle \longrightarrow \alpha_k$$

for all k. We apply this to the elements ψ_m of the chosen orthonormal basis, and we derive

$$\sum_{m \leqslant M} |\langle \psi_m, w_j \rangle|^2 \longrightarrow \sum_{m \leqslant M} |\beta_m|^2$$

for any $M \geqslant 1$, where β_m is α_k for the index such that $\psi_m = \varphi_k$. The left-hand side is the norm of the projection of w_j on the span of $(\psi_1, \ldots, \psi_M)$, and as such is $\leqslant 1$. Hence we get

$$\sum_{m \leqslant M} |\beta_m|^2 \leqslant 1$$

for all $M \geqslant 1$, which shows that the vector

$$v = \sum_{m \geqslant 1} \beta_m \psi_m$$

exists in H, and has norm $\leqslant 1$. But then, since the φ_k are finite linear combinations of the ψ_m, we get

$$\langle w_j, \varphi_k \rangle \longrightarrow \langle v, \varphi_k \rangle$$

as $j \to +\infty$, for all k. Since $\{\varphi_k\}$ is dense in H, it follows (formally) that

$$\langle w_j, w \rangle \longrightarrow \langle v, w \rangle \tag{A.3}$$

as $j \to +\infty$, for all $w \in H$.[3]

We now argue that, due to the compactness of T, this weaker convergence property suffices to ensure that $Tw_j \to Tv$ in H. If this is true, we are done: since $Tw_j = T_{n_j} v_{n_j} + (T - T_{n_j}) v_{n_j}$, we have

$$\|Tw_j\| = \lambda_{n_j} + o(1) \to \lambda,$$

and therefore $\|Tv\| = \lambda$, as desired (note that this also shows that $\|v\| = 1$).

Now, we will finally use the explicit form of T to prove that $Tw_j \to Tv$. We write $k(x, y) = k_x(y)$, so that (using our simplifying assumption that k is continuous) each k_x is a well-defined continuous (hence bounded and square-integrable) function on X. Then we can write

$$(Tw_j - Tv)(x) = \langle w_j - v, \overline{k_x} \rangle,$$

and hence

$$\|Tw_j - Tv\|^2 = \int_X |\langle w_j - v, \overline{k_x} \rangle|^2 d\mu(x).$$

[3] One can think of this property as saying that "every coordinate" of w_j, captured by any linear functional on H, converges to the corresponding coordinate of v. To clarify the meaning of this, note that the formula $\|v\| = \sup_{\|w\| \leqslant 1} |\langle v, w \rangle|$ shows that the convergence in norm is equivalent to a uniform convergence (over w of norm 1) of the corresponding "coordinates".

We can now apply the dominated convergence theorem to conclude: by (A.3), the integrand converges pointwise to 0, and moreover

$$|\langle w_j - v, \overline{k_x}\rangle|^2 \leqslant (\|w_j\| + \|v\|)^2 \|k\|_\infty^2 \leqslant 2\|k\|_\infty^2,$$

which is an integrable function on X, so that the dominated convergence theorem does apply. □

Remark A.2.6. In the language of weak convergence, we can summarize the last steps of the argument as follows: (1) any sequence of unit vectors in H contains a weakly convergent subsequence; (2) a compact operator maps a weakly convergent sequence to a norm-convergent sequence. In general, (1) is a consequence of the Banach-Alaoglu Theorem (see, e.g., [**31**, Th. IV.21]) and the self-duality of Hilbert spaces, while (2) is most easily seen as coming from the alternate characterization of compact operators (on a Hilbert space H) as those mapping the unit ball to a relatively compact subset H (see, e.g., [**31**, VI.5, Th. VI.11] for this fact).

A.3. The Stone-Weierstrass theorem

We conclude with the statement of the general Stone-Weierstrass approximation theorem, which is used in Exercises 5.4.4 and 5.4.6 for an alternative proof of the Peter-Weyl Theorem.

Theorem A.3.1 (Stone-Weierstrass approximation theorem). *Let X be a compact topological space, and let $A \subset C(X)$ be a subspace of continuous functions on X such that* (1) *A is an algebra: if f, $g \in A$, the product fg is also in A;* (2) *A is stable under complex conjugation: if $f \in A$, then $\bar{f}$ is in A (here $\bar{f}$ maps x to $\overline{f(x)}$);* (3) *A separates points, i.e., for any elements x, y in X with $x \neq y$, there exists some $f \in A$ such that $f(x) \neq f(y)$.*

Then A is dense in $C(X)$ for the uniform topology, i.e., for any $f \in C(X)$ and $\varepsilon > 0$, there exists $g \in A$ with

$$\sup_{x \in X} |f(x) - g(x)| < \varepsilon.$$

Bibliography

[1] B. Bekka, P. de la Harpe and A. Valette: *Kazhdan's Property* (T), New Math. Monographs 11, Cambridge Univ. Press (2008).

[2] N. Berry, A. Dubickas, N. Elkies, B. Poonen and C. J. Smyth: *The conjugate dimension of algebraic numbers*, Quart. J. Math. 55 (2004), 237–252.

[3] E.D. Bolker: *The spinor spanner*, American Math. Monthly 80 (1973), 977–984.

[4] A. Borel: *Linear algebraic groups*, 2nd edition, GTM 126, Springer 1991.

[5] W. Bosma, J. Cannon and C. Playoust, *The Magma algebra system, I. The user language* J. Symbolic Comput. 24 (1997), 235–265; also `http://magma.maths.usyd.edu.au/magma/`

[6] A.E. Brouwer and M. Popoviciu: *The invariants of the binary nonic*, J. Symb. Computation 45 (2010) 709–720.

[7] D. Bump: *Automorphic forms and representations*, Cambridge Studies in Adv. Math. 55, Cambridge Univ. Press (1998).

[8] C.W. Curtis and I. Reiner: *Representation theory of finite groups and associative algebras*, AMS Chelsea Publishing, 1962.

[9] P. Diaconis: *Group representations in probability and statistics*, Inst. Math. Stat. Lecture Notes 11, Institute of Math. Statistics (1988).

[10] N. Dunford and J.T. Schwarz: *Linear operators, part II: spectral theory*, Wiley (1963).

[11] P. Etingof, O. Golberg, S. Hensel, T. Liu, A. Schwendner, D. Vaintrob, E. Yudovina: *Introduction to representation theory*, `arXiv:0901.0827`.

[12] G. Folland: *Real Analysis*, Wiley (1984).

[13] W. Fulton and J. Harris: *Representation theory, a first course*, Universitext, Springer (1991).

[14] The GAP Group: *GAP – Groups, Algorithms, and Programming, Version 4.4.9*, 2007; also `www.gap-system.org`

[15] K. Girstmair: *Linear relations between roots of polynomials*, Acta Arithmetica 89 (1999), 53–96.

[16] W.T. Gowers: *Quasirandom groups*, Comb. Probab. Comp. 17 (2008), 363–387.

[17] W.T. Gowers (editor): *The Princeton companion to mathematics*, Princeton Univ. Press (2008).

[18] R. Hartshorne: *Algebraic Geometry*, Grad. Texts in Math. 52, Springer Verlag (1977).

[19] I.M. Isaacs: *Character theory of finite groups*, Academic Press (1976).

[20] I.M. Isaacs: *Finite group theory*, Graduate Studies in Math. 92, American Math. Soc. (2008).

[21] F. Jouve, E. Kowalski and D. Zywina: *An explicit integral polynomial whose splitting field has Galois group* $W(E_8)$, Journal de Théorie des Nombres de Bordeaux 20 (2008), 761–782.

[22] N.M. Katz: *Larsen's alternative, moments and the monodromy of Lefschetz pencils*, in "Contributions to automorphic forms, geometry, and number theory (collection in honor of J. Shalika's 60th birthday)", J. Hopkins Univ. Press (2004), 521–560.

[23] N.M. Katz: *Convolution and equidistribution: Sato-Tate theorems for finite field Mellin transforms*, Annals of Math. Studies, Princeton Univ. Press (to appear).

[24] A.W. Knapp: *Representation theory of semisimple groups*, Princeton Landmarks in Math., Princeton Univ. Press (2001).

[25] E. Kowalski: *The large sieve, monodromy, and zeta functions of algebraic curves, II: independence of the zeros*, International Math. Res. Notices 2008, `doi:10.1093/imrn/rnn091`.

[26] S. Lang: *Algebra*, 3rd edition, Grad. Texts in Math. 211, Springer (2002).

[27] M. Larsen: *The normal distribution as a limit of generalized Sato-Tate measures*, preprint (`http://mlarsen.math.indiana.edu/~larsen/papers/gauss.pdf`).

[28] M.W. Liebeck, E.A. O'Brien, A. Shalev and P.H. Tiep: *The Ore conjecture*, J. Eur. Math. Soc. (JEMS) 12 (2010), 939–1008.

[29] W. Magnus: *Residually finite groups*, Bull. Amer. Math. Soc. 75 (1969), 305–316.

[30] N. Nikolov and L. Pyber: *Product decompositions of quasirandom groups and a Jordan-type theorem*, J. European Math. Soc., to appear.

[31] M. Reed and B. Simon: *Methods of modern mathematical physics, I: Functional analysis*, Academic Press 1980.

[32] M. Reed and B. Simon: *Methods of modern mathematical physics, II: Self-adjointness and Fourier theoretic techniques*, Academic Press 1980.

[33] J. Rotman: *An introduction to the theory of groups*, 4th ed., Grad. Texts in Math. 148, Springer (1995).

[34] J.-P. Serre: *Linear representations of finite groups*, Grad. Texts in Math. 42, Springer (1977).

[35] J.-P. Serre: *Cours d'arithmétique*, Presses Univ. France (1988).

[36] J.-P. Serre: *Moursund lectures*, `arXiv:math.0305257`

[37] J.-P. Serre:*Topics in Galois theory*, Res. Notes in Math., Jones and Bartlett (1992).

[38] S.F. Singer: *Linearity, symmetry, and prediction in the hydrogen atom*, Undergraduate texts in mathematics, Springer (2005).

[39] T.A. Springer: *Invariant theory*, Springer Lecture Notes in Math. 585, (1977).

[40] T.A. Springer: *Linear algebraic groups*, 2nd edition, Progr. Math. 9, Birkhaüser (1998).

[41] S. Sternberg: *Group theory and physics*, Cambridge Univ. Press (1994).

[42] L. Takhtajan: *Quantum mechanics for mathematicians*, A.M.S Grad. Studies in Math. 95, 2008.

[43] N.J. Vilenkin: *Special functions and the theory of group representations*, Translations of Math. Monographs 22, A.M.S (1968).

[44] H. Weyl: *The theory of groups and quantum mechanics*, Dover (1950).

[45] A. Zygmund: *Trigonometric series*, 3rd Edition, Cambridge Math. Library, Cambridge Univ. Press (2002).