

Introduction and motivation

This book is intended to provide a basic introduction to some of the fundamental ideas and results of *representation theory*. In this 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 prerequisites—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 form) is perhaps best summarized by P. Sarnak’s advice in the *Princeton Companion to Mathematics* [24, 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).

Finally, although I have tried to illustrate many aspects of representation theory, there remains many topics that are barely mentioned or omitted altogether. Maybe the most important are:

- The representation theory of anything else than groups; in particular, Lie algebras and their representations only make passing appearances, and correspondingly those aspects of representation theory that really depend on these techniques are not developed in any detail. Here, the book [20] by Fulton and Harris is an outstanding resource, and the book [18] by Etingof, Golberg, Hensel, Liu, Schwendner, Vaintrob, and Yudovina illustrates different aspects, such as the representations of quivers.
- In a related direction, since it really depends on Lie algebraic methods, the precise classification of representations of compact Lie groups, through the theory of highest weight representations, is not considered beyond the case of $SU_2(\mathbf{C})$; this is however covered in great detail in many other texts, such as [20] again, the book [37] of Knapp (especially Chapter V), or the book [35] of Kirillov.

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 to “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.

Thanks to the reviewers for the original book proposal for suggestions and comments—in particular for some well-deserved critical comments concerning certain of the choices of notation in the first version of the text, and for pointing out that Proposition 2.3.23 is false over non-algebraically closed fields.

Finally, many thanks to E. Dunne for reading the whole manuscript carefully and making many suggestions and corrections!

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 one to imagine how widespread and influential the concepts of representation theory turn out to be in current mathematics. Few fields of 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.

The next chapter starts by presenting the fundamental vocabulary that is the foundation of representation theory and by illustrating it with examples. In Chapter 3, we then present a number of short sections concerning variants of the definition of representations: 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. One can also replace groups by other objects: we will mention associative algebras and Lie algebras. These variants are all important topics in their own right, but some will only reappear briefly in the rest of the book.

Continuing, 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 can also serve as a “blueprint” to many generalizations: 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 in Chapter 5. Then Chapter 6 presents some concrete examples of applications involving compact Lie groups (compact matrix groups, such as unitary groups $U_n(\mathbf{C})$)—the most important being perhaps the way representation theory explains a lot about the way the most basic atom, hydrogen, behaves in the real world. . . .

The final Chapter 7 has again a survey flavor, and it is intended to serve as an introduction to two other important classes of groups: algebraic groups, on the one hand, and non-compact locally compact groups, on the other hand. This last case is illustrated through the fundamental example of the group $SL_2(\mathbf{R})$ of two-by-two real matrices with determinant 1. We use it primarily to illustrate some of the striking new phenomena that arise when compactness is missing.

In Appendix A, we have gathered statements and sketches of proofs for certain facts, especially the Spectral Theorem for compact self-adjoint linear operators, which are needed for rigorous treatments of unitary representations of topological groups.

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 discuss again 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 prime numbers in an arithmetic progression, whenever this is not clearly impossible:

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

$$p \equiv a \pmod{q},$$

i.e., such that p is of the form $p = nq + a$ for some $n \geq 1$.

For instance, taking $q = 10^k$ to be a power of 10, we can say that, for whichever ending pattern of digits $\mathbf{d} = d_{k-1}d_{k-2} \cdots d_0$ we might choose, with $d_i \in \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$, *provided* the last digit d_0 is not one of $\{0, 2, 4, 5, 6, 8\}$, there exist infinitely many prime numbers p with a decimal expansion where \mathbf{d} are the final digits. To illustrate this, taking $q = 1000$, $\mathbf{d} = 237$, we find

1237, 2237, 5237, 7237, 8237, 19237, 25237, 26237, 31237, 32237,
38237, 40237, 43237, 46237, 47237, 52237, 56237, 58237, 64237,
70237, 71237, 73237, 77237, 82237, 85237, 88237, 89237, 91237, 92237

to be those prime numbers ending with 237 which are ≤ 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 that 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 non-relativistic theory, there are four quantum numbers for a given pure state of hydrogen, 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

$$n \geq 1, \quad 0 \leq \ell \leq n - 1, \quad -\ell \leq m \leq \ell, \quad s \in \{-1/2, 1/2\}.$$

It is rather striking that much of this quantum mechanical model of the hydrogen atom can be “explained” qualitatively by an analysis of the representation theory of the underlying symmetry group (see [64] or [58]) leading in particular to a natural explanation of the intricate structure of these four quantum numbers! We will attempt to explain the easiest part of this story, which only involves the magnetic and angular momentum quantum numbers, 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 square matrices of size 2 with determinant 1, and with coefficients in the finite field $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z}$. This group is generated by the two elements

$$(1.1) \quad s_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad s_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

(this is a fairly easy fact from elementary group theory, see, e.g., [51, Th. 8.8] for $K = \mathbf{F}_p$ or Exercise 4.6.20). 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 exist an integer $k \geq 1$ and elements g_1, \dots, g_k , each of which belongs to S , such that

$$g = g_1 \cdots g_k.$$

Given g , let $\ell(g)$ be the *smallest* k for which such a representation exists. One may ask, how large can $\ell(g)$ be when g varies over $\mathrm{SL}_2(\mathbf{F}_p)$? The following result gives an answer:

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

$$\ell(g) \leq C \log p$$

for all $g \in \mathrm{SL}_2(\mathbf{F}_p)$.

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.

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

$$g = g_1 \cdots g_k.$$

For instance, what would you do with

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

(for $p \geq 3$)? 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 is 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 [23] (and Nikolov and Pyber [47]) which are much in the same spirit, and we use the same crucial ingredient concerning representations of $\mathrm{SL}_2(\mathbf{F}_p)$. The book [13] of Davidoff, Sarnak, and Valette gives a complete elementary proof and is fully accessible to readers of this book.

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 for the groups \mathbf{R}/\mathbf{Z} and \mathbf{R} , respectively (see Section 7.3). The next motivating example is purely algebraic.

Example 1.2.6 (Burnside's $p^a q^b$ theorem). Recall that a group G is called *solvable* if there is an increasing sequence of subgroups

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

each normal in the next (but not necessarily in G), such that 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 symmetric group \mathfrak{S}_5 of order $5! = 120$ is *not* solvable, and since 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 was only in the late 1960s that a proof not using representation theory was found, first by Goldschmidt when the primes p and q are odd, and then independently by Bender and Matsuyama for the general case. There is a full account of this in [29, §7D], and although it is not altogether overwhelming in length, the reader who

compares them 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 later chapters, other structures are involved: groups are considered with a topology, measure spaces and integration theory is involved, as well as basic Hilbert space theory and functional analysis. All these are used at the level of introductory graduate courses.

We will use the following notation:

(1) For a set X , $|X| \in [0, +\infty]$ denotes its cardinality, with $|X| = \infty$ if X is infinite. There is no distinction in this text between the various infinite cardinals.

(2) We denote by $\mathbf{R}^{+, \times}$ the interval $]0, +\infty[$ seen as a subgroup of the multiplicative group \mathbf{R}^\times .

(3) If k is a field and $d \geq 1$ an integer, an element of $\mathrm{GL}_d(k)$ (or of $\mathrm{GL}(E)$ where E is a finite-dimensional k -vector space) is called *unipotent* if there exists $n \geq 1$ such that $(u - \mathrm{Id}_k)^n = 0$.

(4) 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 .

(5) 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

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

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

(7) A *projection* $f : M \rightarrow M$ is a linear map such that $f \circ f = f$. If f is such a projection, we have $M = \mathrm{Im}(f) \oplus \mathrm{Ker}(f)$; we also say that f is the projection on $\mathrm{Im}(f)$ with kernel $\mathrm{Ker}(f)$.

(8) 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

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

In particular, there is a bilinear map

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

that corresponds 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 a *unique* expression

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

with $v_j \in M$ for all j .

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

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

the sequence is said to be *exact* if $\mathrm{Im}(f) = \mathrm{Ker}(g)$ in M . In particular, a sequence

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

is exact if and only if $\mathrm{Ker}(f) = 0$, which means that f is injective, and a sequence

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

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

A sequence

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

where all three intermediate 3-term sequences are exact, is called a *short exact sequence*. This means that f is injective, g is surjective and the image of f coincides with the kernel of g . It is also usual to say that M is an *extension* of M'' by M' . Note that there is no typo here: this is indeed the standard terminology, instead of speaking of extensions of M' .

(10) Given a vector space E over a field k and a family $(F_i)_{i \in I}$ of linear subspaces of E , we say that the subspaces F_i are *in direct sum* if the subspace they span is a direct sum of the F_i , or in other words, if

$$F_i \cap \left(\sum_{\substack{j \in I \\ j \neq i}} F_j \right) = 0$$

for all $i \in I$ (equivalently, any family $(f_i)_{i \in I}$ of vectors in F_i , which are zero for all but finitely many indices i , is linearly independent).

(11) Given a group G , we denote by $[G, G]$ the *commutator group* (or *derived subgroup*) 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; see Remark 4.4.5 for examples. The subgroup $[G, G]$ is normal in G , and the quotient group $G/[G, G]$ is abelian; it is called the *abelianization* of G .

(12) 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 $q = p^n$, $n \geq 1$, is a power of p . In Chapter 4, we need some simple facts about these, in particular the fact that for each $n \geq 1$, 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). The group homomorphism

$$N = N_{k/\mathbf{F}_p} : \begin{cases} k^\times & \longrightarrow & \mathbf{F}_p^\times \\ x & \longmapsto & \prod_{j=0}^{n-1} x^{p^j} \end{cases}$$

(called the *norm* from k to \mathbf{F}_p) is well defined and surjective. Indeed, it is well defined because one checks that $N(x)^p = N(x)$, and surjective, e.g., because the kernel is defined by a non-zero polynomial equation of degree at most $1 + p + p + \cdots + p^{n-1} = (p^n - 1)/(p - 1)$, and hence contains at most that many 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 $y \in k^\times$.

Similarly, the homomorphism of abelian groups

$$\text{Tr} = \text{Tr}_{k/\mathbf{F}_p} : \begin{cases} \mathbf{F}_q & \longrightarrow & \mathbf{F}_p \\ x & \longmapsto & x + x^p + \cdots + x^{p^{n-1}} \end{cases}$$

is well defined and is surjective; it is called the *trace* from k to \mathbf{F}_p .

(13) 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.

(14) When considering a Hilbert space H , we speak synonymously of an inner product or of a positive-definite hermitian form, which we denote $\langle \cdot, \cdot \rangle$ 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

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

for two vectors v, w and a scalar $\alpha \in \mathbf{C}$. We recall that a Hilbert space is *separable* if it has a finite or countable orthonormal basis. If $T : H_1 \rightarrow H_2$ is a continuous (synonymously, bounded) linear operator between Hilbert spaces, the *adjoint* of T is the unique linear operator $T^* : H_2 \rightarrow H_1$ such that

$$\langle T(v_1), v_2 \rangle_{H_2} = \langle v_1, T^*(v_2) \rangle_{H_1}$$

for all $v_1 \in H_1$ and $v_2 \in H_2$. The operator T is called *self-adjoint* if and only if $T^* = T$ and *unitary* if and only if $TT^* = T^*T = \text{Id}$.

(15) We will always consider Hausdorff topological spaces, except if explicitly mentioned otherwise (this will only happen in Section 7.1).

(16) 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 , and which satisfies the regularity conditions

$$\begin{aligned} \mu(A) &= \inf\{\mu(U) \mid U \supset A, U \text{ open}\} \quad \text{for all Borel sets } A, \\ \mu(U) &= \sup\{\mu(K) \mid K \subset U, K \text{ compact}\} \quad \text{for all open sets } U \end{aligned}$$

(see, e.g., [19, §7.1]); if X is σ -compact (for instance, if X is a separable metric space), then in fact these regularity conditions are automatically satisfied (see, e.g., [19, Th. 7.8]).

(17) The support of a Borel measure μ is the set in X defined as the complement of the union of all open sets U in X such that $\mu(U) = 0$. This definition is useful if either X has a countable basis of open sets, for instance $X = \mathbf{R}$, or if μ is a Radon measure, since in those cases the support of μ is closed; see, e.g., [19, Exercise 2, p. 208].

(18) 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:

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

(19) If $\varphi : X \rightarrow Y$ is a measurable map between two measure spaces and μ is a measure on X , then the *image measure* $\nu = \varphi_*\mu$ on Y is defined by

$$\nu(B) = \mu(\varphi^{-1}(B))$$

for $B \subset Y$ measurable, or equivalently by the integration formula

$$\int_Y f(y) d\nu(y) = \int_X f(\varphi(x)) d\mu(x)$$

for any $f : Y \rightarrow \mathbf{C}$ which is integrable (or measurable and ≥ 0).

(20) Finally, a probability measure μ on an arbitrary measure space X is a measure such that $\mu(X) = 1$; the measure $\mu(A)$ of a measurable subset $A \subset X$ is then also called the *probability* of A .

The language of representation theory

2.1. Basic language

We begin by restating formally the following 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 *dimension of ϱ* , or sometimes its *degree* or *rank*. We will denote it $\dim(\varrho)$.

Remark 2.1.2. It is also customary to just say that ϱ is a k -representation of G , and to omit mentioning the field k if it is clear from context. Similarly, when the homomorphism ϱ is clear from context, one may say only that “ E is a representation of G ”. Another common alternative notation is “let (ϱ, E) be a k -representation of G ”.

Given a representation $\varrho : G \longrightarrow \mathrm{GL}(E)$ 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 sometimes called *G -translates of v* (or simply *translates of v* , when the context is clear). Similarly, when ϱ is clearly understood, one may simply write

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

and these notations are all frequently used.

The basic rules that ϱ satisfies are then the relations

$$(2.1) \quad \varrho(1)v = v \quad (gh)v = \varrho(gh)v = \varrho(g)(\varrho(h)v) = g(hv),$$

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

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

$$\left\{ \begin{array}{ll} G \times E & \longrightarrow E \\ (g, v) & \longmapsto g \cdot v. \end{array} \right.$$

It should be clear already 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 (“tautological” 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 a transformation $\pi_G(g)$ through its action on the basis vectors: we define

$$(2.3) \quad \pi_G(g)e_h = e_{gh}$$

for all $g \in G$ and all basis vectors e_h . Then to check that π_G 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 $\pi_G(1)$ 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 $\pi_G(g_1g_2)$ is $e_{g_1g_2h}$ by definition. And since $\pi_G(g_2)e_h$ is the basis vector e_{g_2h} , we also have

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

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

Another easily defined representation is the *right-regular representation*, or simply *regular* representation ϱ_G of G over k : let $C_k(G)$ be the space¹ of

¹ The notation is not completely standard.

all functions

$$f : G \rightarrow k$$

(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 $\varrho_G(g)$ acting on $C_k(G)$ by the rule

$$\varrho_G(g)f(x) = f(xg)$$

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

$$\varrho_G(gh)f(x) = f(xgh),$$

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

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

which completes the check that $\varrho_G(gh) = \varrho_G(g)\varrho_G(h)$.

Exercise 2.1.3. (1) Show that the formula $\lambda_G(g)f(x) = f(g^{-1}x)$ also defines a representation of G on $C_k(G)$. It is called the *left-regular representation* λ_G of G (over k).

(2) Show that the formula

$$\varrho(g, h)f(x) = f(g^{-1}xh)$$

defines a representation ϱ of $G \times G$ on $C_k(G)$.

In the previous examples, the representation map ϱ is injective (it is clear in the second case 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 $1 \in k^\times = \text{GL}(k)$. This is not injective unless $G = 1$. Note that one should not 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 2.1.4 (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 $\text{Ker}(\varrho) = \{1\}$ 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 $\text{Ker}(\varrho) = G$.

Remark 2.1.5. 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 one-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

$$(2.4) \quad \chi_s : x \mapsto e^{sx}$$

is a one-dimensional representation.

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

$$(2.5) \quad e_m : x \mapsto e^{2i\pi mx}$$

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 $q \geq 1$ 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

$$(2.6) \quad x \mapsto e^{2i\pi mx/q}$$

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 2.1.6 (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 *intertwining operator*, or *intertwiner*, 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

$$\begin{array}{ccc} E_1 & \xrightarrow{\Phi} & E_2 \\ \varrho_1(g) \downarrow & & \downarrow \varrho_2(g) \\ E_1 & \xrightarrow{\Phi} & E_2 \end{array}$$

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 $\text{Hom}(E_1, E_2)$, which we denote $\text{Hom}_G(\varrho_1, \varrho_2)$. (This vector space may of course reduce to 0.)

The following are simple facts, but they are also of crucial importance:

Proposition 2.1.7 (Functoriality). *Let G be a group, and let k be 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 2.1.8 (The category of representations). In the language of category theory (which we will only use incidentally in remarks in this book), 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 . It is also equivalent to

$$\varrho_1 = \Phi^{-1} \circ \varrho_2 \circ \Phi.$$

As another general 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 2.1.9 (Trivial subrepresentations). Consider the case where F is the space of all vectors $v \in E$ 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 2.1.4. This means that if n is the dimension² 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

$$\left\{ \begin{array}{l} \mathbf{1} \hookrightarrow E \\ t \mapsto tv \end{array} \right.$$

²Which may be finite or infinite.

of representations. This gives a k -linear isomorphism

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

(the reciprocal map sending $\Phi : \mathbf{1} \rightarrow E$ 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 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

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

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 2.1.10 (One-dimensional representations). A one-dimensional k -representation χ of a group G is simply a homomorphism $\chi : G \rightarrow k^\times$ (this is because, for any one-dimensional k -vector space, there is a canonical isomorphism $k^\times \rightarrow \text{GL}(V)$, obtained by mapping $\lambda \in k^\times$ to λId). Generalizing Example 2.1.9, which corresponds to $\chi = 1$, for an arbitrary k -representation $\varrho : G \rightarrow \text{GL}(V)$, a non-zero intertwiner $\chi \xrightarrow{\Phi} \varrho$ corresponds to the data of a non-zero vector $v \in V$ such that

$$\varrho(g)v = \chi(g)v$$

for all $g \in G$ (the reader should check this elementary fact). This means exactly that v is a *common eigenvector* for all operators $\varrho(g)$. For instance, χ itself, if seen as a k -valued function on G , is an element of $C_k(G)$ which corresponds to an intertwiner $\chi \rightarrow C_k(G)$, as the reader should check.

Example 2.1.11 (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 Example 2.1.9. A very useful fact is that if we take the vectors invariant under a *normal subgroup* of G , we still obtain a subrepresentation of E , though usually not a trivial one.

Lemma 2.1.12 (Invariants under normal subgroups). *Let k be a field, let G be a group, and let $H \triangleleft G$ be 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 $w = \varrho(g)v \in E^H$. 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. \square

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 2.1.13 (Regular representation). Consider the two examples of representations π_G and ϱ_G associated to a group G and field k that were discussed just after the Definition 2.1.1. We claim that π_G (acting on $k(G)$) is isomorphic to a subrepresentation of ϱ_G (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} ; in other words

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

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 $g \cdot \Phi(e_h) = \varrho_G(g)\Phi(e_h)$ maps x to

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

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 $\text{Im}(\Phi)$ is always equal to the subspace of functions which are zero except at finitely many points.

Remark 2.1.14. 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 last is of less importance in this book, but we will mention it briefly nevertheless, and it is very important in certain areas of number theory. The other two are, however, of fundamental importance.

2.2. Formalism: changing the space

This part of 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 is 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 that the reader 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 and, 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 $E' = \text{Hom}_k(E, k)$, and the associated representation of G acting on the space of k -linear maps $\text{End}_k(E) = \text{Hom}_k(E, E)$.

As will be seen, only the last one may not be 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 let G be a group. Let T be any covariant 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

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

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

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

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 2.1.7.

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.

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 E/F also has a natural linear action of G , simply induced by ϱ : given $v \in E/F$ and $g \in G$, the action $g \cdot v$ is the image in E/F of $\varrho(g)\tilde{v}$ for any $\tilde{v} \in E$ mapping to v under the canonical surjective map $E \rightarrow E/F$. 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 E/F .

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

$$E \longrightarrow E/F$$

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

In the same vein, given 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 $\text{Ker}(\Phi) \subset E_1$ is a subrepresentation of E_1 ;
- the image $\text{Im}(\Phi) \subset E_2$ is a subrepresentation of E_2 ;
- the natural linear isomorphism

$$E_1/\text{Ker}(\Phi) \simeq \text{Im}(\Phi)$$

(induced by Φ) is an isomorphism of representations;

- The cokernel $\text{Coker}(\Phi) = E_2/\text{Im}(\Phi)$ 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 2.1.9, 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

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

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

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

This condition is equivalent to $\ker(\lambda) \supset E_1$, where E_1 is the subspace of E spanned by all vectors of the form

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

or equivalently it corresponds to a linear form

$$E/E_1 \longrightarrow k$$

extended to E by composition $E \longrightarrow E/E_1 \longrightarrow k$.

Note that E_1 is also a subrepresentation of E , since

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

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 $g \cdot v \equiv v$ 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 . It is the *largest* quotient of ϱ that is a trivial representation of G (like the invariant space, it may well be zero) and by the above, we can write

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

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

Exercise 2.2.2. Show that if $H \triangleleft G$, the H -coinvariant space E_H has an induced structure of representation of G . (This is the analogue, for the coinvariants, of Lemma 2.1.12.)

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

$$(2.10) \quad (\varrho_1 \oplus \varrho_2)/\varrho_1 \simeq \varrho_2,$$

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

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

where, since we do not distinguish between infinite cardinals, the convention that the dimension is infinite if either there is an infinite-dimensional ϱ_i , or

if I is infinite and infinitely many summands ϱ_i are non-zero. Equally useful are the natural isomorphisms

$$\begin{aligned}\mathrm{Hom}_G(\varrho, \varrho_1 \oplus \varrho_2) &\simeq \mathrm{Hom}_G(\varrho, \varrho_1) \oplus \mathrm{Hom}_G(\varrho, \varrho_2), \\ \mathrm{Hom}_G(\varrho_1 \oplus \varrho_2, \varrho) &\simeq \mathrm{Hom}_G(\varrho_1, \varrho) \oplus \mathrm{Hom}_G(\varrho_2, \varrho),\end{aligned}$$

and similarly for an arbitrary (finite) number of summands.³

Exercise 2.2.3. Let G be a group, let k be a field, and let ϱ_1, ϱ_2 be two k -representations of G , acting on E_1 and E_2 , respectively. Show that a k -linear map $\Phi : E_1 \rightarrow E_2$ is a G -homomorphism if and only if the graph

$$\Gamma = \{(v, \Phi(v)) \mid v \in E_1\} \subset E_1 \oplus E_2$$

of Φ is a subrepresentation of $\varrho_1 \oplus \varrho_2$.

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

$$(2.11) \quad \varrho/\varrho_1 = \varrho/\mathrm{Im}(\Phi) \simeq \varrho_2$$

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

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

as k -vector spaces, *it is not always the case that E_2 can be found as a subrepresentation of ϱ* . When the complement E_2 can be chosen to be a subrepresentation of E , this subrepresentation (say π_2) is necessarily isomorphic to ϱ_2 (since $\pi_2 \simeq (\varrho_1 \oplus \pi_2)/\varrho_1 \simeq \varrho/\varrho_1 \simeq \varrho_2$, as representations of G).

A useful equivalent criterion for the existence of such a complementary subrepresentation is the following:

Lemma 2.2.4. *Let G be a group, let k be a vector space, and let $\varrho : G \rightarrow \mathrm{GL}(E)$ be a representation. Let $E_1 \subset E$ be a subrepresentation of E .*

(1) *Suppose $E_2 \subset E$ is a subspace of E complementary to E_1 , so that $E = E_1 \oplus E_2$. Then E_2 is a subrepresentation of E if and only if the linear projection map*

$$\Phi \begin{cases} E & \longrightarrow E \\ v_1 + v_2 & \longmapsto v_1 \end{cases}, \quad v_1 \in E_1, v_2 \in E_2,$$

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

³ 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.

(2) *There exists a linear complement E_2 which is a subrepresentation if and only if there exists an idempotent $\Phi \in \text{Hom}_G(E, E)$ which is a projection, such that $\text{Im}(\Phi) = E_1$. In this case, $E_2 = \ker \Phi$ is such a complement.*

Proof. Assertion (1) is elementary, and (2) is of course a consequence of (1). This follows by noting first that if Φ is an idempotent, the kernel $\text{Ker } \Phi = E_2$ 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 $\rho(g)v = \rho(g)v_1 + \rho(g)v_2$ with $\rho(g)v_i \in E_i$ again, and hence $\Phi(\rho(g)v) = \rho(g)v_2 = \rho(g)\Phi(v)$. \square

In certain circumstances, a subrepresentation complementary to ρ_1 always exists (for instance, for finite groups when k has characteristic 0, as we will discuss in Chapter 4). Here is a standard example where it fails. Consider the additive group $G = \mathbf{R}$, the field $k = \mathbf{R}$, and the representation

$$(2.12) \quad \rho \begin{cases} G & \longrightarrow & \text{GL}_2(\mathbf{R}) \\ x & \longmapsto & \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}. \end{cases}$$

(We leave it as an exercise to check, if needed, that this is a homomorphism.) In terms of the canonical basis (e_1, e_2) of \mathbf{R}^2 , this means that

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

The subspace $E_1 = \mathbf{R}e_1$ is a subrepresentation of \mathbf{R}^2 , 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* subrepresentation E_2 which is complementary to E_1 . This can be checked either by a direct computation or more abstractly. For the former approach, note that E_2 would be of dimension 1. Now let $f = \alpha e_1 + \beta e_2$ be any vector in \mathbf{R}^2 , and assume that $\rho(x)f \in \mathbf{R}f$ for all $x \in \mathbf{R}$. We take $x = 1$ and deduce that there exists $\lambda \in \mathbf{R}$ such that

$$\rho(1)f = (\alpha + \beta)e_1 + \beta e_2 = \lambda f = \lambda\alpha e_1 + \lambda\beta e_2.$$

This can only happen if $\beta = 0$, in which case $f \in \mathbf{R}e_1$, so that the line spanned by f is in fact the same as E_1 .

The more abstract argument runs as follows: the quotient representation \mathbf{R}^2/E_1 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 $\rho(x)$ are all equal to 1). Thus, if E_2 were to exist, we would have, by the above, an isomorphism

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

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

We come back to the general case where (2.11) holds. Given intertwiners

$$\varrho_1 \longrightarrow \varrho_2 \longrightarrow \varrho_3$$

of k -representations of a group G , one says that the sequence is an *exact sequence (of representations)* if and only if it is exact as a sequence of k -vector spaces. Similarly one can speak of *short exact sequences* of representations, and (by linear algebra) the situation (2.11) can be summarized by a *short exact sequence*

$$0 \rightarrow \varrho_1 \longrightarrow \varrho \xrightarrow{\Phi} \varrho_2 \rightarrow 0 ;$$

again using the terminology for modules, one also says that ϱ is an *extension* of ϱ_2 by ϱ_1 .

When there is a morphism $\varrho_2 \xrightarrow{\Psi} \varrho$ such that $\Phi \circ \Psi = \text{Id}$, one says that the exact sequence *splits*. This happens precisely when the space of ϱ contains a subrepresentation complementary to ϱ_1 (necessarily isomorphic to ϱ_2), so that $\varrho \simeq \varrho_1 \oplus \varrho_2$. 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 2.2.5 (Irreducible, semisimple, isotypic representations). Let G be a group, and let k be 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 (some of the ϱ_i may be isomorphic.)

(3) A semisimple k -representation ϱ of G is *isotypic* if it is a direct sum of irreducible subrepresentations that are all isomorphic; if these subrepresentations are all isomorphic to a representation π , then one says that ϱ is π -isotypic.

We will see later that, up to permutation, the irreducible summands of a semisimple representation 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 2.2.6, is that these “building blocks” are “incommensurable”, in some sense: two non-isomorphic irreducible representations can have “no interaction”.

Lemma 2.2.6 (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 k -representation ϱ of G , any G -homomorphism $\pi \rightarrow \varrho$ is either 0 or injective, and any G -homomorphism $\varrho \rightarrow \pi$ is either 0 or surjective.*

(2) *Given irreducible k -representations π and ϱ of G , a homomorphism $\pi \rightarrow \varrho$ 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.) \square

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 2.2.7 (Existence of irreducible subrepresentations). *Let G be a group, let k be a field, and let ϱ be 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 2.2.8 (Cyclic vector). It is tempting to suggest a more general argument by saying that, given a non-zero representation $G \rightarrow \mathrm{GL}(E)$ 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 a prime number and the representation on \mathbf{C}^2 given by

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

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

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

and since

$$\begin{vmatrix} 1 & 1 \\ 1 & e^{2i\pi/p} \end{vmatrix} = e^{2i\pi/p} - 1 \neq 0,$$

this vector does *generate* the whole space.

For a given representation $\varrho : G \rightarrow \mathrm{GL}(E)$ 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

$$\varrho = \bigoplus_{1 \leq i \leq k} \varrho_i,$$

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 2.7.28 below) that v is a cyclic vector for ϱ .

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

$$\chi : G \rightarrow \mathrm{GL}(E),$$

where E is a one-dimensional vector space over k . As we noted already, since $\mathrm{GL}(E)$ is *canonically* isomorphic to k^\times by the homomorphism mapping $\lambda \in k^\times$ to $k \mathrm{Id}$, this homomorphism χ is just a homomorphism $\chi : G \rightarrow k^\times$.

Such homomorphisms are sometimes called *characters* of G , although this clashes with the more general notion of character that we will see below in Definition 2.7.36. We use this terminology in the next easy proposition:

Proposition 2.2.9. *Let G be a group, and let k be a field.*

- (1) *A character $G \rightarrow k^\times$ is irreducible.*

(2) Two characters $\chi_1, \chi_2 : G \longrightarrow k^\times$ are isomorphic if and only if they are equal as functions on G .

We give the argument but urge the reader to try to check this if the result seems unclear.

Proof. (1) A one-dimensional vector space contains no non-zero proper subspace at all, and must therefore be irreducible.

(2) Since $\mathrm{GL}(E) \simeq \mathrm{GL}(k) = k^\times$ for any one-dimensional space E , we can assume that χ_1 and χ_2 both act by scalar multiplication on k . Then an intertwiner $\Phi : k \longrightarrow k$ is given by $\Phi = \lambda \mathrm{Id}$ for some fixed $\lambda \in k$, and it is an isomorphism if and only if $\lambda \neq 0$. The intertwining condition becomes

$$\Phi(\chi_1(g)x) = \chi_2(g)\Phi(x),$$

for all $x \in k$, which implies $\lambda\chi_1(g) = \lambda\chi_2(g)$ for all $g \in G$. Clearly, this is possible with $\lambda \neq 0$ if and only if $\chi_1 = \chi_2$. \square

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

$$E \simeq \bigoplus_{i \in I} ke_i$$

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 two-dimensional representation in (2.12) is not semisimple. Indeed, since we saw that it is not irreducible, it can only be semisimple if the space \mathbf{R}^2 decomposes as a direct sum $\mathbf{R}^2 = \mathbf{R}f_1 \oplus \mathbf{R}f_2$, where each summand is a subrepresentation, but we checked above that $\mathbf{R}e_1$ is the only one-dimensional subrepresentation of ρ .⁴

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

Lemma 2.2.10 (Stability of semisimplicity). *Let G be a group, and let k be a field.*

(1) *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.*

⁴ Note that we will soon be able to use a more abstract argument: from the Jordan–Hölder–Noether Theorem below, one sees that if ρ were semisimple, it would be trivial, which it is not.

(2) If $\varrho : G \longrightarrow \mathrm{GL}(E)$ is an arbitrary representation of G and if E_1, E_2 are semisimple subrepresentations of E , then their sum $E_1 + E_2 \subset E$, whether it is a direct sum or not, is 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 2.2.11 (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 $F_1 \subset E$ of ϱ , there exists a complementary subrepresentation, i.e., a G -stable subspace $F_2 \subset E$ such that

$$E = F_1 \oplus F_2.$$

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 2.2.10. (1) Let ϱ act on E , and let $F \subset E$ be a subrepresentation. We are going to check that the condition of Lemma 2.2.11 applies to F .⁵ Thus, let $F_1 \subset F$ be a subrepresentation of F ; it is also one of E , hence there exists a subrepresentation $F_2 \subset E$ such that

$$E = F_1 \oplus F_2.$$

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

Now we claim that $F = F_1 \oplus (F \cap F_2)$, which shows that F_1 has also a stable complement $F \cap F_2$ in F , and finishes the proof that F is semisimple. Indeed, F_1 and $(F \cap F_2)$ are certainly in direct sum, and if $v \in F$ and we write $v = v_1 + v_2$ with $v_1 \in F_1$, $v_2 \in F_2$, we also obtain

$$v_2 = v - v_1 \in F \cap F_2$$

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

(2) The sum $E_1 + E_2$ of two subrepresentations of E is isomorphic to a quotient of the direct sum $E_1 \oplus E_2$ by the surjective linear map

$$\begin{cases} E_1 \oplus E_2 & \longrightarrow & E_1 + E_2 \\ (v_1, v_2) & \longmapsto & v_1 + v_2. \end{cases}$$

Since this map is an intertwiner, and the direct sum of semisimple representations is semisimple (which is easy to see from the definition), it follows from (1) that $E_1 + E_2$ is also semisimple. \square

Proof of Lemma 2.2.11. 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 will proceed to check this. However, we first check that \tilde{F} exists. This is easy if E is finite dimensional, and in general it follows from a quite standard application of Zorn's Lemma (see, e.g., [40, p. 693]). We give the details, which the reader may wish to supply independently.

Consider the set \mathcal{D} of subrepresentations \tilde{F} of E such that $F \cap \tilde{F} = 0$, ordered by inclusion of subspaces. We wish to find a maximal element of this ordered set. The subspace 0 is an element of \mathcal{D} , which is therefore not empty. By Zorn's Lemma, it suffices to show that any totally ordered subset of \mathcal{D} has an upper bound in \mathcal{D} . Let $\mathcal{S} \subset \mathcal{D}$ be such a subset. Consider the subset

$$\tilde{F} = \bigcup_{V \in \mathcal{S}} V \subset E.$$

From the fact that \mathcal{S} is totally ordered, it follows that \tilde{F} is a linear subspace of E , and that it is a subrepresentation (e.g., for the last property, for any $g \in G$ and $v \in \tilde{F}$, there exists $V \in \mathcal{S}$ such that $v \in V$, and then $g \cdot v \in V \subset \tilde{F}$). Obviously, we have $V \subset \tilde{F}$ for all $V \in \mathcal{S}$, and hence \tilde{F} is an upper bound of

\mathcal{S} for inclusion, and it is enough to show that $\tilde{F} \in \mathcal{D}$ to conclude that it is also an upper bound for \mathcal{S} in \mathcal{D} . But if $v \in \tilde{F} \cap F$, there exists $V \in \mathcal{S}$ such that $v \in V \cap F = 0$, by definition of \mathcal{D} , and hence we get the desired result.

Now, given a subrepresentation \tilde{F} obtained by this construction, consider for every i the intersection

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

Since E_i is an irreducible representation of G , this intersection is either 0 or equal to E_i . In fact, it cannot 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 2.2.7, and otherwise it requires some care but can be done, as explained in Exercise 2.2.13 below.

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 $J \subset I$.

First, we again use Zorn's Lemma to show that there exists a maximal subset $J \subset I$ (with respect to inclusion) such that the sum of the E_i , $i \in J$, is a direct sum. Indeed, we say that a subset $J \subset I$ is *direct* if the E_i , $i \in J$, are in direct sum. We order the set \mathcal{D} of direct subsets of I by inclusion. Any singleton $\{i\}$ is direct, and if \mathcal{J} is any totally ordered subset of \mathcal{D} , the union $K = \bigcup_{J \in \mathcal{J}} J$ is an upper bound for \mathcal{D} for inclusion. We leave to the reader the exercise of checking that K is also direct, and hence is an upper bound for \mathcal{J} in \mathcal{D} , which allows to apply Zorn's Lemma to \mathcal{D} .

Now fix a maximal direct set J , and let F be the direct sum of those E_i , $i \in J$. For any $i \notin J$, the intersection $E_i \cap F$ cannot be zero, as this would allow us to replace J by $J \cup \{i\}$, 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. \square

Remark 2.2.12. 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

$$E = F \oplus F'$$

for some irreducible subspace F and complementary representation F' . The proof of Lemma 2.2.10 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 2.2.13 (Existence of irreducible subrepresentation). Let G be a group, let k be a field, and let $\varrho : G \rightarrow \mathrm{GL}(E)$ be a completely reducible k -representation of G , with $E \neq 0$. We want to show that such a representation 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

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

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 tensor (but we recall that $E_1 \otimes E_2$ is not simply the set of such pure tensors, although they generate the tensor product as k -vector space).

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 are clear from context. For the same type of general reasons, all the standard isomorphisms between tensor products, such as

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

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 finitely many factors, which is independent of the order of the product.⁶

Similarly, we have

$$\varrho \otimes \left(\bigoplus_i \varrho_i \right) \simeq \bigoplus_i (\varrho \otimes \varrho_i).$$

⁶ In the theory of automorphic representations, an important role is played by certain special infinite tensor products; see [11] or [3].

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 $\dim(\varrho_1 \otimes \varrho_2) = (\dim \varrho_1)(\dim \varrho_2)$. In particular, if χ_1 and χ_2 both have dimension 1, so does the tensor product, and in fact since χ_1 and χ_2 take values in k^\times , the tensor product $\chi_1 \otimes \chi_2$ is just the product of functions $g \mapsto \chi_1(g)\chi_2(g) \in k^\times$. It is customary to omit the tensor product in the notation in that case, writing just $\chi_1\chi_2$.

Exercise 2.2.14 (Tensor product by a one-dimensional representation). Let $\varrho : G \rightarrow \mathrm{GL}(E)$ be a k -representation of a group G , and let $\chi : G \rightarrow \mathrm{GL}(k) = k^\times$ be a one-dimensional representation.

(1) Show that $\varrho \otimes \chi$ is isomorphic to the representation

$$g \mapsto \chi(g)\varrho(g)$$

of G on the same space E . One sometimes says that $\varrho \otimes \chi$ is obtained by *twisting* ϱ by χ .

(2) Show that $\varrho \otimes \chi$ is irreducible (resp. semisimple) if and only if ϱ is irreducible (resp. semisimple).

(3) Show that $\varrho \otimes (\chi_1\chi_2) \simeq (\varrho \otimes \chi_1) \otimes \chi_2$ for any two one-dimensional representations of G .

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 $\varrho : G \rightarrow \mathrm{GL}(E)$ is a k -representation of G , we can construct the following:

- the symmetric powers $\mathrm{Sym}^m(E)$ of E , for $m \geq 0$;
- the alternating powers $\bigwedge^m E$, for $m \geq 0$.

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

$$G \rightarrow \mathrm{GL}(\mathrm{Sym}^m(E)), \quad G \rightarrow \mathrm{GL}(\bigwedge^m E),$$

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

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

From elementary multilinear algebra, we recall that if E has finite dimension, the full symmetric algebra $\text{Sym}(E)$ is infinite dimensional, but the alternating algebra is not; indeed, $\bigwedge^m E = 0$ if $m > \dim E$. More generally, the dimensions of the symmetric and alternating powers are given by

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

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

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

Remark 2.2.15 (Symmetric powers as coinvariants). Let E be a k -vector space. For any $m \geq 1$, 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

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

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

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

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

Exercise 2.2.16. Let G be a group, let k be a field, and let $\varrho : G \rightarrow \text{GL}(E)$ be a finite-dimensional k -representation of G of dimension $d = \dim(E)$.

(1) Let $\delta = \bigwedge^d \varrho$ be the d -th alternating power of ϱ . Show that δ is a one-dimensional k -representation of G isomorphic to the representation $g \mapsto \det(\varrho(g))$.

(2) Let $F \subset E$ be a linear subspace. Show that F is a subrepresentation of E if and only if the one-dimensional subspace $\bigwedge^{\dim(F)} F$ is a subrepresentation of $\bigwedge^{\dim(F)} E$. (*Hint:* Recall that vectors (v_1, \dots, v_k) in E are linearly independent if and only if $v_1 \wedge \cdots \wedge v_k$ is non-zero in $\bigwedge^k E$.)

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 $E' = \text{Hom}_k(E, k)$, which is called the *contragredient* $\check{\varrho}$ of ϱ . More precisely, since the transpose reverses the order of composition,⁷ 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 $\check{\varrho}(g)\lambda$ is the linear form

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

Remark 2.2.17. One way to remember this is to write the definition in the form of the equivalent invariance formula

$$(2.13) \quad \langle g \cdot \lambda, g \cdot v \rangle = \langle \lambda, v \rangle$$

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

In fact, this definition turns out to be mostly useful when E is finite dimensional, because it is associated to the *algebraic* dual space of E . When considering topological groups and representations satisfying continuity assumptions, another dual representation can be defined using the topological dual, and it has better properties (see Section 3.3, in particular Lemma 3.3.7).

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 proposition shows how the contragredient interacts with some of the other operations previously discussed.

Proposition 2.2.18. *Let k be a field, G a group.*

(1) *The contragredient is functorial: given k -representations ϱ_1 and ϱ_2 of G acting on E_1 and E_2 , respectively, and an intertwiner $\Phi : \varrho_1 \rightarrow \varrho_2$, the transpose ${}^t\Phi$ is an intertwiner $\check{\varrho}_1 \rightarrow \check{\varrho}_2$.*

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

$$\left(\bigoplus_i \varrho_i\right)^\vee \simeq \bigoplus_i \check{\varrho}_i.$$

⁷ Equivalently, in the language of Proposition 2.2.1, the assignment $T(E) = E'$ is a contravariant functor which *reverses* arrows in contrast with (2.8), i.e., $T(f \circ g) = T(g) \circ T(f)$, with $T(f)$ the transpose.

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

$$(\varrho_1 \otimes \varrho_2)^\vee \simeq \check{\varrho}_1 \otimes \check{\varrho}_2.$$

(4) 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 \check{F} of its contragredient, given by

$$\begin{aligned} F &\mapsto F^\perp = \{\lambda \in E' \mid \lambda(F) = 0\}, \\ \check{F} &\mapsto {}^\perp\check{F} = \{x \in E \mid \lambda(x) = 0 \text{ for all } \lambda \in \check{F}\}. \end{aligned}$$

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

(6) If ϱ is one dimensional, then the contragredient of ϱ is the one-dimensional representation given by $\check{\varrho}(g) = \varrho(g^{-1}) = \varrho(g)^{-1}$.

Recall that the transpose in (1) is defined by

$$\langle {}^t\Phi(\lambda), v \rangle = \langle \lambda, \Phi(v) \rangle$$

for $\lambda \in E'_2$, $v \in E_1$.

Proof. Almost all these properties are elementary consequences of linear algebra. We only give some details concerning (4) and (6). For the former, we observe that both constructions indicated send, in any case, a subrepresentation of ϱ to one of $\check{\varrho}$ or conversely. This follows by the defining formula

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

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 $\check{\varrho}$) 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 $\check{\varrho}$ is.

Without the finite-dimension assumption, we can still argue in this manner to show that if $\check{\varrho}$ is irreducible, the original representation is also. For any subrepresentation F of ϱ , the subrepresentation F^\perp of $\check{\varrho}$ must be either 0 or all of $\check{\varrho}$. 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.

Finally, for (6), if ϱ is one dimensional, then so is the contragredient $\check{\varrho}$. If we view ϱ as acting by multiplication on k , then the linear form defined

by $\lambda(x) = x$ (for $x \in k$) is a basis of k' , and we find that

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

i.e., that $\check{\varrho}(g)\lambda = \varrho(g^{-1})\lambda$, which gives the result.⁸ □

Remark 2.2.19. The absence of symmetry in some parts of this lemma is not surprising because dual spaces of infinite-dimensional vector spaces do not behave very well in the absence of topological restrictions (see, e.g., [9, §7, no. 5, th. 6].)

Exercise 2.2.20 (Contragredient and invariants). Let G be a group, let k be a field, and let $\varrho : G \longrightarrow \mathrm{GL}(E)$ be a k -representation of G .

(1) Show that there exists a natural isomorphism

$$(E')^G \simeq (E_G)'$$

of k -vector spaces, where the left-hand side is the space of invariants of the contragredient of ϱ acting on E' and the right-hand side is the dual of the coinvariant space of E (see Section 2.2.2).

(2) If $\dim(E) < +\infty$, show that there exists a natural isomorphism

$$(E^G)' \simeq (E')_G,$$

i.e., the dual of the invariant space is isomorphic to the coinvariants of the contragredient. (*Hint:* Use (1) and duality.)

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

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

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 linear map

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

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

This isomorphism shows that, if

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

are k -representations of G , the endomorphism space $\mathrm{Hom}_k(E, F)$ carries a natural representation of G . Indeed, we can define an action simply by asking that the isomorphism (2.14) be an isomorphism of representations.

⁸ Note that it is only because a one-dimensional representation takes values in the abelian group k^\times that $g \mapsto \varrho(g^{-1})$ is a homomorphism!

It is useful to have a more direct description of this action, and this leads to a definition which does not require the representations to be finite dimensional. We state this as a proposition:

Proposition 2.2.21 (Action on homomorphisms). *Let k be a field, let G be a group, and let*

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

be k -representations of G . Then G acts on $\mathrm{Hom}_k(E, F)$ by

$$(2.15) \quad (g \cdot \Phi) = \tau(g)\Phi\varrho(g)^{-1} : E \longrightarrow F$$

for $g \in G$ and $\Phi : E \longrightarrow F$.

If π denotes this representation and if $\dim(\tau) < +\infty$, then we have an isomorphism

$$(2.16) \quad \pi \simeq \check{\varrho} \otimes \tau.$$

Furthermore, for any ϱ and τ , we have

$$(2.17) \quad \mathrm{Hom}_k(\varrho, \tau)^G = \mathrm{Hom}_G(\varrho, \tau),$$

the space of intertwiners between ϱ and τ .

Note that the definition of the representation on $\mathrm{Hom}_k(E, F)$ is such that the diagrams

$$\begin{array}{ccc} E & \xrightarrow{\Phi} & F \\ \varrho(g) \downarrow & & \downarrow \tau(g) \\ E & \xrightarrow{\pi(g)\Phi} & F \end{array}$$

commute for all $g \in G$. Concretely, we thus have

$$(2.18) \quad (g \cdot \Phi)(v) = g \cdot \Phi(g^{-1} \cdot v)$$

for all $v \in E$. Yet another way to remember this is to write the formula in the form

$$(2.19) \quad (g \cdot \Phi)(g \cdot w) = g \cdot \Phi(w)$$

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

Proof. We leave it to the reader to check that (2.15) defines a representation of G . If we grant this, we note that the important relation (2.17) is an immediate consequence of the definition.

We now check the isomorphism (2.16), which is in fact the same as (2.14). This means that we must check that this linear isomorphism is an intertwiner between π and $\check{\varrho} \otimes \tau$. This 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 let $A = A_{\lambda, v}$ be the associated homomorphism. Then the rank 1 map associated to

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

is given by

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

This property exactly states that the linear isomorphism (2.14) is an isomorphism of representations. \square

Remark 2.2.22. The simplest example of the action (2.15) comes from the natural representation of $\mathrm{GL}_n(k)$ on k^n . It is then the same as the action of $\mathrm{GL}_n(k)$ on the space $M_n(k)$ of square matrices of size n by conjugation: $g \cdot A = gAg^{-1}$ for any $g \in \mathrm{GL}_n(k)$ and any matrix A .

These representations on homomorphism spaces are extremely useful and are used in many contexts to compare two representations. This arises from the relation (2.17) which identifies the space $\mathrm{Hom}_G(\varrho_1, \varrho_2)$ of G -homomorphisms between ϱ_1 and ϱ_2 with the invariant space in $\mathrm{Hom}_k(\varrho_1, \varrho_2)$. This interpretation makes it possible to understand and study intertwining operators from within representation theory. For instance, from one of the parts of Schur's Lemma 2.2.6, we see that $\mathrm{Hom}_k(\pi, \varrho)^G = 0$ if π and ϱ are non-isomorphic irreducible representations. We suggest looking at the proof of Proposition 2.8.2 below for another good illustration of the use of the homomorphism representation to compare two representations.

Remark 2.2.23 (Other actions 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

$$(2.20) \quad (\tau(g)A)(w) = \varrho_2(g)(A(w)),$$

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 shorthand such as $(g \cdot A)w = g \cdot Aw$, and which is also less clear if $\varrho_1 = \varrho_2$, and we consider representations on $\mathrm{End}_k(\varrho_1)$).

Exercise 2.2.24. (1) Show that the representation τ just described is isomorphic to a direct sum of $\dim E$ copies of ϱ_2 (i.e., to a direct sum of $\dim E$ representations, each of which is isomorphic to ϱ_2 ; in particular, if ϱ_2 is irreducible, the representation τ is ϱ_2 -isotypic). (*Hint:* For a basis $(w_j)_{j \in J}$ of

E , show that the map

$$(2.21) \quad \bigoplus_{j \in J} \varrho_2 \longrightarrow \tau$$

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

$$A(w_j) = v_j$$

is an isomorphism.)

(2) Define similarly a representation $\tilde{\tau}$ on $\text{Hom}_k(E, F)$ by putting

$$(\tilde{\tau}(g)A)(w) = A(\varrho_1(g^{-1})w).$$

Show that this is a k -representation of G . Show that if $\dim(\varrho_1)$ is finite, then $\tilde{\tau}$ is isomorphic to a direct sum of $\dim(F)$ copies of $\check{\varrho}_1$.

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

$$\varrho : G \rightarrow \text{GL}(E)$$

of G , simply by composition

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

The underlying vector space is therefore unchanged, 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 the language of category theory, 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 $\text{Res}_H^G(\varrho)$, 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 $\phi : G \rightarrow H$ 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, let K be a normal subgroup, and let $H = G/K$ 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 $K \subset \text{Ker}(\pi)$.

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

There is a very important special case of this discussion. Recall that the derived subgroup $G' = [G, G]$ of G , generated by commutators, has the property that for any group G and abelian group A , there is a canonical bijection between homomorphisms $G \rightarrow A$ and homomorphisms $G/[G, G] \rightarrow A$. Since $\text{GL}_1(k) = k^\times$ is abelian, applying this to $A = k^\times$, we obtain

Proposition 2.3.3 (One-dimensional representations). *Let G be a group, and let $G^{ab} = G/[G, G]$ be the abelianization of G . For any field k , the one-dimensional representations of G correspond with the homomorphisms*

$$G^{ab} \longrightarrow k^\times.$$

In particular, if G is perfect, i.e., if $[G, G] = G$, then 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 $d \geq 2$ and k is any field, $\text{SL}_d(k)$ is known to be perfect except when $d = 2$ and $k = \mathbf{F}_2$ or $k = \mathbf{F}_3$ (see, e.g., [40, Th. 8.3, Th. 9.2]). Thus, no such group has a non-trivial one-dimensional representation.

Exercise 2.3.4 (Semisimplicity of restriction). Let G be a group, let k be a field, and let $\varrho : G \rightarrow \text{GL}(E)$ be a semisimple k -representation of G . Let $H \subset G$ be a finite-index normal subgroup of G . Show that $\text{Res}_H^G(\varrho)$ is also semisimple as a representation of H . (*Hint*: One can assume that ϱ is

irreducible—show that there exists a maximal semisimple subrepresentation of $\text{Res}_H^G(\varrho)$.)

The converse of this statement is not true without restrictions, but it is valid when k has characteristic 0; see Exercise 4.1.2.

One of the most basic and important constructions of representation theory, and in some sense the first notion that may not be immediately related to notions of linear algebra,⁹ is the operation of *induction*. We will now define it and spend a fair amount of time discussing its basic properties, and it will reappear throughout the book.

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 about possible ways of doing this. We also recommend reading what follows first by assuming ϕ to be an inclusion map, and removing it from the notation.

One defines the induced¹⁰ representation as follows: given

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

we define first the k -vector space

$$(2.22) \quad F = \{f : G \rightarrow E \mid f(\phi(h)x) = \varrho(h)f(x) \text{ for all } h \in H, x \in G\}$$

(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 π of G (namely the restriction π to F of the analogue of the regular representation) defined by

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

for $f \in F$, $g \in G$, and $x \in G$. Indeed, we need only check that F is stable under the regular representation of G , which is true, because F is defined using conditions relating to multiplication on the left by elements of H . Formally, if $f_1 = \pi(g)f$, we find that

$$f_1(\phi(h)x) = f(\phi(h)xg) = \varrho(h)f(xg) = \varrho(h)f_1(x),$$

⁹ It is, however, related to certain tensor products.

¹⁰ It is unfortunate that the terminology “induced” may clash with the use of the adjective “induced” in less formal senses, and that “induction” conflicts with, e.g., proofs by induction.

for all $h \in H$ and $x \in G$, which means that—as desired—we have $f_1 \in F$ again.

Especially when ϕ is an inclusion, one writes

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

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 ϕ explicitly). One may even drop H and G from the notation when they are clear from the context.

Remark 2.3.5. If we take $h \in \text{Ker}(\phi)$, the transformation formula in (2.22) for elements of F gives

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

so that, in fact, any function $f \in F$ takes values in the space $E^{\text{Ker}(\phi)}$ of invariants of E under the action of the subgroup $\text{Ker}(\phi)$ 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 2.3.11 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 2.1.12, this subspace $E^{\text{Ker}(\phi)}$ is in fact a subrepresentation of E , so that in the condition

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

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

Example 2.3.6 (Elementary examples of induction). (1) By the definition of F and comparison with the definition of the regular representation, we see that the latter can be expressed as

$$(2.23) \quad C_k(G) = \text{Ind}_1^G(\mathbf{1}),$$

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

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

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

for any k -representation $\varrho : G \rightarrow \text{GL}(E)$ of G , the map $F \rightarrow E$ giving this isomorphism being simply

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

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

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

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

we then claim that we have

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

with the action of G/K induced by ϱ (note that by Lemma 2.1.12, the subspace E^K is a subrepresentation of E). This isomorphism is again given by $f \mapsto f(1)$, which—as we have remarked—is a vector in $E^{\mathrm{Ker}(\phi)} = E^K$, as the reader is again invited to check.

(4) Suppose now that $\phi : G \rightarrow G$ 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

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

satisfies

$$\Phi(g \cdot f) = f(g) = f(\varrho(\varrho^{-1}(g))) = \varrho(\varrho^{-1}(g))f(1) = \phi_*\varrho(f),$$

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 $x \in G$), then 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 $g \in G$, so that it is an isomorphism $\phi_*\varrho \longrightarrow \varrho$.

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

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

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

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

where the first map is a quotient map, the second is the isomorphism induced by ϕ , and the third is an injection. (This is also a special case of *induction in stages*; see Proposition 2.3.20 below.)

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

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

In that case, the space F of $\text{Ind}_H^G(\varrho)$ 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 :

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

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 $\text{GL}_2(\mathbf{F}_q)$, where \mathbf{F}_q is a finite field.

Remark 2.3.7. 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 those 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 requires some changes to behave reasonably. The model (2.22) is then a good definition as it can immediately suggest that we consider restricted classes of functions on G instead of all of them (see Example 5.2.10 and Section 7.4.)

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

Proposition 2.3.8 (Functoriality of induced representations). *Let k be a field, and let $\phi : H \rightarrow G$ be a group homomorphism. For any ho-*

homomorphism $\varrho_1 \xrightarrow{\Phi} \varrho_2$ of k -representations of H , there is a corresponding homomorphism

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

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

Proposition 2.3.9 (Frobenius reciprocity; adjointness of induction and restriction). *Let k be a field, and let $\phi : H \rightarrow G$ be a group homomorphism. For any k -representation ϱ_1 of G and ϱ_2 of H , there is a natural isomorphism of k -vector spaces*

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

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

The last isomorphism, or its immediate corollary

$$(2.25) \quad \dim \text{Hom}_G(\varrho_1, \text{Ind}_H^G(\varrho_2)) = \dim \text{Hom}_H(\text{Res}_H^G(\varrho_1), \varrho_2)$$

is called the *Frobenius reciprocity formula*. As we will see many times, it is an extremely important result. In fact, in some (precise) sense, it characterizes the induced representation, and can almost be said to define it (see Remark 2.3.21 for an explanation). We will use induction and Frobenius reciprocity extensively to analyze the properties and the decomposition of induced representations.

We also remark that the definition of the induced representation that we chose is the best for handling situations in which $[G : H]$ can be infinite. If $[G : H]$ is finite, there is another natural model (say $\underline{\text{Ind}}_H^G$) which leads to isomorphisms

$$(2.26) \quad \text{Hom}_G(\underline{\text{Ind}}_H^G(\varrho_1), \varrho_2) \simeq \text{Hom}_H(\varrho_1, \text{Res}_H^G(\varrho_2)),$$

and those are sometimes considered to be the incarnation of Frobenius reciprocity (see Exercise 2.3.16 and, e.g., [28, Ch. 5]).

Proof of Proposition 2.3.8. The induced homomorphism $\Phi_* = \text{Ind}(\Phi)$ is easy to define using the model of the induced representation given above: denoting by F_1, F_2 the spaces corresponding to $\text{Ind}_H^G(\varrho_1)$ and $\text{Ind}_H^G(\varrho_2)$, respectively, we define $\Phi_*(f)$ for $f \in F_1$ to be given by

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

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

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

for all $h \in H$ shows that $\Phi_*(f)$ 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 . \square

Proof of Proposition 2.3.9. Here also 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 as defined in (2.22).

We will describe how to associate to

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

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

$$(2.27) \quad T(\Phi) : \begin{cases} F_1 & \longrightarrow & E_2 \\ v & \mapsto & \Phi(v)(1). \end{cases}$$

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

$$\Psi : F_1 \rightarrow E_2,$$

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

$$(2.28) \quad x \mapsto \Psi(\varrho_1(x)v)$$

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 now finish checking that these constructions give the Frobenius reciprocity isomorphisms. First, we check that $f = \tilde{\Psi}(v)$ is indeed in F_2 : for all $x \in G$ and $h \in H$, we have

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

(using the fact that Ψ is a homomorphism from $\text{Res}_H^G(\varrho_1)$ to ϱ_2). Next, $\tilde{\Psi}$ intertwines ϱ_1 and $\text{Ind}_H^G(\varrho_2)$: for $g \in G$, the function $\tilde{\Psi}(\varrho_1(g)v)$ is

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

and this coincides with

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

The remaining property we need is that the two constructions are inverse of each other. If we start with $\Psi \in \text{Hom}_H(F_1, E_2)$, then construct $\tilde{\Psi} = \tilde{T}\Psi$, the definitions (2.27) and (2.28) 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 $\Phi \in \text{Hom}_G(F_1, F_2)$, 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

$$\Phi(\varrho_1(x)v) = x \cdot \Phi(v) = (y \mapsto \Phi(v)(yx)).$$

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

Example 2.3.10. 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 2.3.11. *Let k be a field, and let $\phi : H \rightarrow G$ be a group homomorphism. For a k -representation ϱ of H , acting on a space E , we have*

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

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

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

Remark 2.3.12. Note that this formula is one case where one must be careful in the infinite-dimensional case. We mentioned in Section 1.3 that we do not distinguish between infinite cardinals, and with this convention the formula is valid (i.e., both sides are infinite if and only if one of them is). However, the formula is *not* true if one interprets the left-hand side as the cardinal (say c_1) of a basis of $\mathrm{Ind}_H^G(\varrho)$ and the right-hand side as the product (say c_2) of the cardinals of $[G : \mathrm{Im}(\phi)]$ and that of a basis of $E^{\mathrm{ker} \phi}$, if the index $[G : \mathrm{Im}(\phi)]$ is infinite (and $E^{\mathrm{ker} \phi}$ is non-zero): we then have $c_1 > c_2$.

Proof. The idea is very simple: the definition of the space F on which the induced representation acts shows that the value of $f \in F$ 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 f ; each seems to belong to the space E , but as we observed in Remark 2.3.5, it is in fact constrained to lie in the possibly smaller space $E^{\mathrm{Ker}(\phi)}$.

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

$$\tilde{f} : R \rightarrow E^{\mathrm{Ker}(\phi)},$$

and we consider the obvious k -linear map

$$F \rightarrow \tilde{F}$$

defined by restricting functions on G to R (using Remark 2.3.5 to see that this is well defined). Now we claim that this is an isomorphism of vector spaces. This implies the formula for the dimension of F (distinguishing the infinite-dimensional case from the finite-dimensional one).

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

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

for all $x \in R$ and $h \in H$. Since these elements, by definition, cover all of G , we get $f = 0$. 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 $\tilde{f} : R \rightarrow E^{\text{Ker}(\phi)}$, we then *define*

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

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

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

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

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

since κ acts trivially on the space $E^{\text{Ker}(\phi)}$, and (as in Lemma 2.1.12) the vector

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

does belong to it. We are now done because

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

finishes the proof that $f \in F$. □

Remark 2.3.13. (1) From the proof we see that one could have defined the induced representation as the k -vector space of all functions

$$\text{Im}(\phi) \backslash G \longrightarrow E^{\text{Ker}(\phi)}$$

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

(2) See Exercise 4.2.8 for an application to proving a lower bound for the minimal index of a proper subgroup of a finite group.

Exercise 2.3.14. Let G be a group, let $H \subset G$ be a subgroup, and let $\varrho : G \rightarrow \text{GL}(E)$ be a k -representation of H .

(1) Show that if F is a subrepresentation of E , then $\text{Ind}_H^G(F)$ is naturally isomorphic to a subrepresentation of $\text{Ind}_H^G(\varrho)$.

(2) Show that if $\text{Ind}_H^G(\varrho)$ is irreducible, then so is ϱ . Is the converse true?

The proof of Proposition 2.3.11 implicitly reveals more information than the dimension of the induced representation. In particular, we can use it to give one answer to the question of recognizing when a given representation of a group G is induced from a subgroup. Not only is this useful in practice (see Proposition 2.8.1, for instance), but it certainly helps in visualizing what the operation of induction is.

Proposition 2.3.15. *Let G be a group, and let k be a field. Let $\varrho : G \rightarrow \text{GL}(E)$ be a finite-dimensional k -representation of G . Assume that there exists a finite-index subgroup $H \subset G$ and a direct sum decomposition*

$$E = \bigoplus_{i=1}^{[G:H]} E_i,$$

where each $E_i \subset E$ is H -stable, such that for any i , we have $E_i = \varrho(g_i^{-1})E_1$ for some $g_i \in G$. Then the representation ϱ is isomorphic to the induced representation $\text{Ind}_H^G(E_1)$.

Proof. Note that the assumption implies that the dimension of E is $[G : H] \dim E_1$, which is the dimension of $\text{Ind}_H^G(E_1)$, as we have just seen, so the result is certainly plausible. We will construct an intertwining map

$$\Phi : \text{Ind}_H^G(E_1) \rightarrow E$$

and show that it is injective. Since we also assume that $\dim E < +\infty$, this will be enough to finish the proof.

The definition of Φ is easy. For f in the space F of the induced representation $\text{Ind}_H^G(E_1)$, we define

$$\Phi(f) = \sum_{i=1}^{[G:H]} \varrho(g_i^{-1})f(g_i).$$

This defines a k -linear map to E , since $f(g_i) \in E_1$ by definition of F . In fact the assumption shows that $\varrho(g_i^{-1})f(g_i) \in E_i$ for each i , and since these spaces are in direct sum, we also immediately see that

$$\text{Ker } \Phi = \{f \in F \mid f(g_i) = 0 \text{ for all } i\}.$$

To deduce that Φ is injective, and indeed also to show that Φ is an intertwiner, we now claim that the (g_i) form a set of representatives of left H -cosets in G , i.e., that

$$G = \bigcup_i Hg_i$$

with the union being disjoint. The number of g_i is the right one, so it suffices to prove that the g_i are in distinct H -cosets. But if $hg_i = g_j$ for some $h \in H$, we deduce that

$$\varrho(g_i^{-1})\varrho(h)E_1 = \varrho(g_j^{-1})E_1,$$

and since E_1 is H -stable, this means that $E_i = E_j$, which means that $i = j$.

We first apply this to prove injectivity of Φ : $f \in \text{Ker } \Phi$ means that f is zero on a set of representatives of $H \setminus G$, and (as in the proof of Proposition 2.3.11), this means that $f = 0$.

We conclude by checking that Φ is an intertwiner. Consider $g \in G$ and $f \in F$. Then the definition gives

$$\Phi(g \cdot f) = \sum_i \varrho(g_i^{-1})f(g_i g).$$

Multiplication on the right by g permutes the left H -cosets: there exists a permutation σ of the indices such that for each i , we have

$$(2.29) \quad g_i g = h_{\sigma(i)} g_{\sigma(i)}$$

for some $h_{\sigma(i)} \in H$. Since $f \in F$, we deduce that

$$\Phi(g \cdot f) = \sum_i \varrho(g_i^{-1})f(g_i g) = \sum_i \varrho(g_i^{-1} h_{\sigma(i)})f(g_{\sigma(i)}).$$

But (2.29) gives

$$g g_{\sigma(i)}^{-1} = g_i^{-1} h_{\sigma(i)}$$

so that, rearranging the sum, this becomes

$$\Phi(g \cdot f) = \sum_i \varrho(g g_{\sigma(i)}^{-1})f(g_{\sigma(i)}) = \varrho(g)(\Phi(f)),$$

concluding the proof. \square

Exercise 2.3.16. Let G be a group, let $H \subset G$ be a finite-index subgroup, and let k be a field. Let $\{g_1, g_2, \dots, g_k\}$ be a set of representatives for right

H -cosets in G with $g_1 = 1$. For a k -representation $\varrho : H \rightarrow \mathrm{GL}(E)$, define a representation π of G as follows. The space F of π is

$$F = \bigoplus_{i=1}^k g_i E,$$

where $g_i E$ denotes a vector space isomorphic to E , and the action is obtained by formally using the action of H on E (given by ϱ) and the requirement that $g_i E$ is the translate of $g_1 E = E$ by g_i , i.e. if $g \in G$, $v \in E$ and

$$gg_i = g_j h$$

for some j and $h \in H$, we put

$$g \cdot (g_i v) = g_j (\varrho(h)v) \in g_j E.$$

(1) Show that π is indeed a representation of G on F , and that it is isomorphic to $\mathrm{Ind}_H^G(\varrho)$.

(2) Show that for any k -representation $\varrho_2 : G \rightarrow \mathrm{GL}(F_2)$, there exists a canonical isomorphism

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

(this is the alternate formula (2.26) for Frobenius reciprocity).

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 2.3.17. 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 2.3.3 can help.)

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

Proposition 2.3.18 (Projection formula). *Let k be a field, and let $\phi : H \rightarrow G$ be 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 (2.24), the proof is not very difficult as the isomorphism can be described explicitly, but the 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, then try to verify that it satisfies the required properties.¹¹

Proof. We denote by F_1 the space of ϱ_1 , by E_2 that of ϱ_2 , and by F_2 the space (2.22) of the induced representation $\text{Ind}_H^G(\varrho_2)$. Moreover, we denote by τ the representation

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

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

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

also defined using (2.22).

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

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

which extends by linearity the definition

$$\Phi(f \otimes v) = (x \mapsto f(x) \otimes x \cdot v)$$

for $f \in F_2$ and $v \in F_1$.

Note that the right-hand side is indeed a function $G \rightarrow E_2 \otimes F_1$, and that $E_2 \otimes F_1$ is the space of τ (in this proof, we write $x \cdot v$ for the action of ϱ_1 on F_1). It is also a bilinear expression of the arguments f and v . Hence, to see that Φ is well defined, it is enough to check that its image does lie in \tilde{F}_2 . But if $\tilde{f} = \Phi(f \otimes v)$, then (using the fact that $f \in F_2$), we obtain

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

for all $x \in G$, $h \in H$, which is the property required for a function $G \rightarrow E_2 \otimes F_1$ 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 $f \otimes v$. Let $\tilde{f} = \Phi(f \otimes v)$ and $g \in G$; then we have

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

which we can also write as

$$f_1(x) \otimes x \cdot w,$$

where $f_1(x) = f(xg) = g \cdot f(x)$ and $w = g \cdot v$, or in other words as

$$\Phi(f_1 \otimes w)(x) = \Phi(g \cdot (f \otimes v))(x),$$

as desired.

¹¹ 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.

To conclude, it remains 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

$$(2.30) \quad w = \sum_j w_j(x) \otimes (x \cdot v_j)$$

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

$$\sum_j f_j \otimes v_j$$

for some functions $f_j \in F_2$. Let us assume such an element is in $\text{Ker}(\Phi)$. This means that for all $x \in G$, we have

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

Thus by the uniqueness of the representations (2.30), we get

$$f_j(x) = 0$$

for all j , or in other words $f_j = 0$ for all j , and this gives $\text{Ker}(\Phi) = 0$.

We now come to surjectivity. Let $\tilde{f} \in \tilde{F}_2$ be given. Again by the observation above, for any $x \in G$, we can write uniquely¹²

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

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

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

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

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

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

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

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

¹² This is the trick: use (2.30) for a varying x , not for a single fixed basis.

by definition, while the right-hand side is

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

Comparing using the uniqueness of (2.30) with x replaced by $\phi(h)x$, we find that, for all j , we have

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

and this does state that each coefficient function \tilde{f}_j is in F_2 . \square

Remark 2.3.19. If ϕ is an injective homomorphism and the groups G and H are finite, then all spaces involved are finite dimensional. Since Proposition 2.3.11 shows that both sides of the projection formula are of degree $[G : H] \dim(\varrho_1) \dim(\varrho_2)$, the injectivity of Φ is sufficient to finish the proof.

Yet another important property of induction (and restriction) is the following, which is called *induction in stages* in the case of induction:

Proposition 2.3.20 (Transitivity). *Let k be a field, 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

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

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 mostly omitting the homomorphisms ϕ , ϕ_1 , ϕ_2 in the notation.

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

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

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 $f \in F$, a function

from G to E , it is natural to consider

$$g \cdot f = (x \mapsto f(xg)),$$

the image of f under the regular representation on E -valued functions. Then $g \mapsto g \cdot f$ is a function from G to F , so its values are themselves functions from G to E . We want $T(f) \in F_2$, so it must be an F_1 -valued function on G , i.e., $T(f)(g)$ must be a function from H_1 to E . Hence it seems plausible to define

$$T(f)(g) = (g \cdot f) \circ \phi_1.$$

If H_1 is a subgroup of G , this is the restriction to H_1 —in the sense of restricting a function on G to one on H_1 —of the function $g \mapsto g \cdot f$ on G .

We can then check that $T(f)$ is, in fact, F_1 -valued; if we omit the group homomorphisms involved, this amounts to letting $\psi = T(f)$ and writing

$$\psi(h_2h_1) = g \cdot f(h_2h_1) = f(h_2h_1g) = \varrho(h_2)f(h_1g) = \varrho(h_2)\psi(h_1),$$

for $h_i \in H_i$, using of course in the middle the assumption that f is in F . Again, this is unlikely to make much sense until the reader has tried and succeeded independently to follow the computation.

Now we should check that $T(f)$ is not only F_1 -valued, but also lies in F_2 , i.e., transforms under H_1 like the induced representation $\text{Ind}_{H_2}^{H_1}(\varrho)$. 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 .¹³ We now start with $\psi \in F_2$ and must define a function from G to E . Unraveling in two steps, we set

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

($\psi(g)$ 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 $f = \tilde{T}(\psi)$, we again let the reader check that

$$f(h_2g) = \psi(h_2g)(1) = (h_2 \cdot \psi(g))(1) = \psi(g)(h_2) = \varrho(h_2)\psi(g)(1) = \varrho(h_2)f(g)$$

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

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

$$(g \cdot f)(1) = f(g),$$

¹³ 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.

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

$$\begin{aligned}\tilde{\psi}(g)(h_1) &= (g \cdot f)(h_1) = f(h_1g) \\ &= \psi(h_1g)(1) = (h_1 \cdot \psi(g))(1) = \psi(g)(h_1)\end{aligned}$$

(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 (since ψ and $\tilde{\psi}$ are functions on G whose values are functions from H_1 to $E \dots$). Thus T and \tilde{T} are reciprocal isomorphisms. \square

Remark 2.3.21 (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 uses character theory; it is restricted to special situations, and will be discussed later. We sketch the second now, since the reader is presumably well motivated to hear about abstract nonsense if it cuts down on the calculations.

The keyword is the adjective “natural” (or “canonical”) that we attributed to the isomorphisms (2.24) 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 2.3.9 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

$$(\Psi : \pi \rightarrow \varrho) \in V(\pi)$$

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

$$\begin{array}{ccc} \mathrm{Hom}_G(\pi, \varrho_1) & \xrightarrow{I(\pi)} & \mathrm{Hom}_G(\pi, \varrho_2) \\ \downarrow & & \downarrow \\ \mathrm{Hom}_G(\pi', \varrho_1) & \xrightarrow{I(\pi')} & \mathrm{Hom}_G(\pi', \varrho_2) \end{array}$$

commute for any $\Phi : \pi' \rightarrow \pi$, 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 (2.24). 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 $\tilde{\varrho}$ with given homomorphism groups $\mathrm{Hom}_G(\pi, \tilde{\varrho})$ which behave naturally. Precisely, the behavior under morphisms must be compatible: we get

Fact bis. Let $\phi : H \rightarrow G$ 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

$$\begin{array}{ccc} \mathrm{Hom}_G(\pi, \varrho') & \xrightarrow{i(\pi)} & \mathrm{Hom}_H(\mathrm{Res}(\pi), \varrho) \\ \downarrow & & \downarrow \\ \mathrm{Hom}_G(\pi', \varrho') & \xrightarrow{i(\pi')} & \mathrm{Hom}_H(\mathrm{Res}(\pi'), \varrho) \end{array}$$

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 (2.22)) are such that the diagrams

$$(2.31) \quad \begin{array}{ccc} \mathrm{Hom}_G(\pi, \mathrm{Ind}_H^G(\varrho)) & \xrightarrow{i(\pi)} & \mathrm{Hom}_H(\mathrm{Res}_H^G(\pi), \varrho) \\ \downarrow & & \downarrow \\ \mathrm{Hom}_G(\pi', \mathrm{Ind}_H^G(\varrho)) & \xrightarrow{i(\pi')} & \mathrm{Hom}_H(\mathrm{Res}_H^G(\pi'), \varrho) \end{array}$$

commute. These are the same as the ones above, with $\varrho' = \mathrm{Ind}(\varrho)$. This is the real content of the observation that the Frobenius reciprocity isomorphisms are natural. Thus the *construction* of (2.22) proved the *existence* of the induced representation characterized by the abstract property of Frobenius reciprocity.

We can now 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

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

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 be easily convinced (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 (2.31).

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.

We now prove the Fact above, using the notation in that statement. Take $\pi = \varrho_1$, so that we have

$$\mathrm{Hom}_G(\pi, \varrho_1) = \mathrm{Hom}_G(\varrho_1, \varrho_1)$$

and $I(\pi) = I(\varrho_1)$ is an isomorphism $\mathrm{Hom}_G(\varrho_1, \varrho_1) \longrightarrow \mathrm{Hom}_G(\varrho_1, \varrho_2)$.

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

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

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 precomposition with I for any π . Indeed, $I(\pi)$ is an isomorphism

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

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

$$\begin{array}{ccc} \text{Hom}_G(\varrho_1, \varrho_1) & \xrightarrow{I(\varrho_1)} & \text{Hom}_G(\varrho_1, \varrho_2) \\ \downarrow & & \downarrow \\ \text{Hom}_G(\pi, \varrho_1) & \xrightarrow{I(\pi)} & \text{Hom}_G(\pi, \varrho_2). \end{array}$$

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

$$I(\varrho_1)(\text{Id}_{\varrho_1}) \circ \Phi = I \circ \Phi \in \text{Hom}_G(\pi, \varrho_2).$$

But if we follow the down-then-right route, we get $I(\pi)(\text{Id}_{\varrho_1} \circ \Phi) = I(\pi)(\Phi)$, and hence the commutativity of these diagrams says that, for all Φ , we have

$$(2.32) \quad I(\pi)(\Phi) = I \circ \Phi,$$

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

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

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

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

which translates, from the formula (2.32) we have just seen (applied with $\pi = \varrho_2$) to

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

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

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

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 how—especially if the reader does (or has done) all computations independently—part 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 proof of the following lemma to the reader:

Lemma 2.3.22. *Let k be a field, and let $\phi : H \rightarrow G$ be a group homomorphism with $\phi(H)$ of finite index in G . For any finite-dimensional representations ϱ and ϱ_i of H , we have natural isomorphisms*

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

and

$$\mathrm{Ind}_H^G\left(\bigoplus_{i \in I} \varrho_i\right) \simeq \bigoplus_{i \in I} \mathrm{Ind}_H^G(\varrho_i)$$

for I finite.

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

$$[G : H] \dim(\varrho_1) \dim(\varrho_2), \quad [G : H]^2 (\dim \varrho_1) (\dim \varrho_2),$$

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* of ϱ_1 and ϱ_2 ,

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

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 2.3.23 (Irreducible representations of direct products). *Let k be an algebraically closed 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 the end of Section 2.7. The statement is false, in general, over non-algebraically closed fields (see Example 2.7.33).

Remark 2.3.24 (Relation with the ordinary tensor product). Consider a group G . There is an injective *diagonal* homomorphism

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

If ϱ_1 and ϱ_2 are k -representations of G , the definitions show that

$$\text{Res}_G^{G \times G}(\varrho_1 \boxtimes \varrho_2) = \varrho_1 \otimes \varrho_2.$$

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

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

into a representation (which we denote $\varrho \otimes K$) over K , by composing with the group homomorphism

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

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

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

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 .

Sometimes, given a field extension K/k (for instance, with K algebraically closed) and a certain property $\mathcal{P}(\rho)$ of a representation ρ , it may happen that $\mathcal{P}(\rho)$ does not hold for a k -representation ρ , but that $\mathcal{P}(\rho \otimes K)$ does (or conversely). If K is an algebraic closure of k , and if $\rho \otimes K$ has the desired property, one then says that the k -representation ρ “has \mathcal{P} absolutely”.

Example 2.4.1. We give here an example of a representation which is irreducible but not *absolutely irreducible*. Consider the (infinite) abelian group $G = \mathbf{R}/\mathbf{Z}$, and the two-dimensional real representation given by

$$\rho : \begin{cases} G \longrightarrow \mathrm{GL}_2(\mathbf{R}) \\ \theta \mapsto \begin{pmatrix} \cos(2\pi\theta) & \sin(2\pi\theta) \\ -\sin(2\pi\theta) & \cos(2\pi\theta) \end{pmatrix} \end{cases}$$

(which corresponds to the action of \mathbf{R}/\mathbf{Z} on the real plane by the rotation with angle $2\pi\theta$). This makes it clear that this is a homomorphism (which is otherwise easy to check using trigonometric identities), and it also makes it clear that ρ is irreducible (there is no non-zero real subspace of \mathbf{R}^2 which is stable under all rotations $\rho(\theta)$, except \mathbf{R}^2 itself).

However, the irreducibility breaks down when extending the base field to \mathbf{C} . Indeed, on \mathbf{C}^2 , we have

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

and

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

This means 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

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

and the other to

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

(its *complex conjugate*, in the sense of Example 2.4.2 below). Thus 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 $\sigma : k \rightarrow k$, and we define a representation $\rho^{(\sigma)}$ by the composition

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

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 as 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 $\rho(g)$ in the basis (v_i) of E_σ is obtained by applying σ^{-1} to all coefficients of the matrix that represents $\rho(g)$. Indeed, for any i , we can write

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

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

This operation on representations can be interesting because ρ and $\rho^{(\sigma)}$ 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 $\rho^{(\sigma)}$ 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 $G \rightarrow \mathrm{GL}(E)$, where E is a \mathbf{C} -vector space, there is a naturally associated *conjugate* representation $\bar{\rho}$ obtained by applying the construction above to the complex conjugation. If E is finite dimensional, then from the basic theory of characters (Proposition 2.7.38 below and Example 2.7.41) one derives the fact that $\bar{\rho}$ is isomorphic to ρ if and only if the function $g \mapsto \mathrm{Tr} \rho(g)$ is real valued. (This can already be checked when ρ is one-dimensional, since $\bar{\rho}$ is then the conjugate function $G \rightarrow \mathbf{C}$, 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 $\bar{\rho}$ are not isomorphic.

Field extensions (including automorphisms) are the only morphisms for fields. However, there are sometimes other possibilities for changing fields, which are more subtle. Suppose for instance that

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

is a complex representation of degree $d \geq 1$ of some group G , and that *with respect to some chosen basis of E* , the image of ϱ is given by matrices with integral coefficients. If we then fix a prime number p , we may consider the reduction modulo p (say $\bar{\varrho}(g)$) of these matrices, which will be elements in $\mathrm{GL}_d(\mathbf{Z}/p\mathbf{Z})$. The properties of the reduction modulo p imply that $g \mapsto \bar{\varrho}(g)$ is a homomorphism from G to $\mathrm{GL}_d(\mathbf{Z}/p\mathbf{Z})$, or in other words, an \mathbf{F}_p -representation of G .

More abstractly, this definition corresponds to the existence of a G -stable lattice M of E , which is an abelian group $M \subset E$ such that $M \otimes_{\mathbf{Z}} \mathbf{C} \simeq E$ and such that $\varrho(g)m \in M$ for all $g \in G$ and $m \in M$. We can then define an \mathbf{F}_p -representation of G on M/pM , which is a d -dimensional \mathbf{F}_p -vector space, simply because pM is also G -stable (by linearity).

This construction can be extremely useful and important. However, it is delicate: first of all, it is not always defined (the G -stable lattice M may not exist), and also it may not be *well defined*, in the sense that taking a different G -stable lattice (there is no uniqueness, since for instance pM works just as well as M) might lead to a non-isomorphic \mathbf{F}_p -representation of G . We refer to [53, §15.2] for further discussion of this theory.

Exercise 2.4.3. Let $G = \mathbf{Z}/2\mathbf{Z}$, and let ϱ be the two-dimensional regular representation of G on \mathbf{C}^2 , with canonical basis e_1, e_2 . Let $f_1 = e_1 + e_2$, $f_2 = e_1 - e_2$. Let $k = \mathbf{Z}/2\mathbf{Z}$.

Show that $M = \mathbf{Z}^2$ is a G -stable lattice and that the k -representation of G on $M/2M$ is not semisimple. On the other hand, show that $M' = f_1\mathbf{Z} \oplus f_2\mathbf{Z}$ is another G -stable lattice, and that the k -representation of G on $M'/2M'$ is trivial, in particular semisimple.

In Exercise 2.6.6, we will see an example of an irreducible representation that reduces modulo a prime to one which is not.

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 the homomorphism

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

that maps g to the matrix representing $\varrho(g)$ in the chosen basis. Indeed, this is what we already 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 (especially because of the choice of a basis), 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 may concatenate bases (e_1, \dots, e_n) of E_1 and (f_1, \dots, f_m) of E_2 to obtain a basis

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

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

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

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

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

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

$$(f'_1, \dots, f'_m) = (\Phi(f_1), \dots, \Phi(f_m))$$

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

$$(2.33) \quad g \mapsto \varrho^m(g) = \left(\begin{array}{c|c} \varrho_1^m(g) & \star \\ \hline 0 & \varrho_2^m(g) \end{array} \right),$$

where ϱ_1^m is the matrix representation in (e_1, \dots, e_n) and ϱ_2^m the one in (f'_1, \dots, f'_m) . The block denoted \star is an important invariant of the short exact sequence. If we view it as a map

$$c : G \longrightarrow M_{n,m}(k)$$

from G to rectangular $n \times m$ matrices with coefficients in k , then c is not a homomorphism, but writing down the relation

$$\varrho^m(gh) = \varrho^m(g)\varrho^m(h),$$

we see that it satisfies instead

$$c(gh) = \varrho_1^m(g)c(h) + c(g)\varrho_2^m(h)$$

for $g, h \in G$.

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 this basis is ordered as

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

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

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

It should be noted however that this explicit form is very rarely useful.

The matrix representation of the contragredient of a representation ϱ is also easy to describe. We have

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

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, but we will see some examples in the next chapters.

2.6. Examples

We collect here some more examples of representations. The first one, in particular, is very important, and it will reappear frequently in various ways in the rest of the book.

2.6.1. Binary forms and invariants. Let k be any field. For an integer $m \geq 0$, 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

$$X^i Y^{m-i}, \quad 0 \leq i \leq m.$$

In fact, these monomials are independent, and therefore form a 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 multiplication of the row vector (X, Y) by the matrix g . In other words, if

$$(2.34) \quad g = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

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).

We then have

Theorem 2.6.1 (Irreducible representations of SL_2). (1) For $k = \mathbf{C}$, the representations ϱ_m , for $m \geq 0$, are irreducible representations of $\mathrm{SL}_2(\mathbf{C})$. In fact, ϱ_m is an irreducible representation of the subgroup $\mathrm{SU}_2(\mathbf{C}) \subset \mathrm{SL}_2(\mathbf{C})$.

(2) Any finite-dimensional, continuous, irreducible representation of the subgroup $\mathrm{SU}_2(\mathbf{C})$ is isomorphic to one of the ϱ_m .

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

The first assertion will be proved in Example 2.7.11 and Exercise 2.7.13, and the second in Theorem 5.6.3 (the notion of continuous representation is explained in Section 3.3). The analogous statement for $\mathrm{SL}_2(\mathbf{C})$ is more complicated: it is not true that the representations ϱ_m exhaust (up to isomorphism) all finite-dimensional irreducible continuous representations of $\mathrm{SL}_2(\mathbf{C})$. This can be seen quickly by noting that the complex conjugate $\bar{\varrho}_n$ of ϱ_m (as defined in Example 2.4.2), which is irreducible, is not isomorphic to any ϱ_n . Indeed, just by comparing dimensions, the only possibility would be that $\bar{\varrho}_n \simeq \varrho_n$, but from simple computations, one can see that the character $g \mapsto \mathrm{Tr}(\varrho_m(g))$ is not real valued, which is a necessary condition for ϱ_m to be isomorphic to its complex conjugate (see (2.50)). One can show that an irreducible finite-dimensional continuous representation of $\mathrm{SL}_2(\mathbf{C})$ is of the form

$$\varrho_m \otimes \bar{\varrho}_n$$

for some integers $m, n \geq 0$ (see, e.g., [36, end of II.3]).

We can however already explain the statement (3). 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 $\dim V_p = p + 1 \geq 3$) and V_p is a subrepresentation. Indeed, for g given by (2.34), we have

$$\begin{aligned} (g \cdot X^p) &= (aX + cY)^p = a^p X^p + c^p Y^p \in W, \\ (g \cdot Y^p) &= (aX + cY)^p = b^p X^p + d^p Y^p, \end{aligned}$$

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 $1 \leq j \leq p-1$). One can also show that $W \subset V_p$ does not have a stable complementary subspace, so 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 $m \geq 0$. 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.)

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 , and recover the result without using matrix representations.

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 $m \geq n \geq 0$, the tensor product $\varrho_m \otimes \varrho_n$ is semisimple¹⁴ and decomposes as*

$$(2.35) \quad \varrho_m \otimes \varrho_n \simeq \varrho_{m+n} \oplus \varrho_{m+n-2} \oplus \cdots \oplus \varrho_{m-n}.$$

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 tautological inclusion $\mathrm{SL}_2(\mathbf{C}) \rightarrow \mathrm{GL}_2(\mathbf{C})$), we see that all other representations ϱ_m arise by taking tensor products iteratively and decomposing them, e.g.,

$$\varrho_1 \otimes \varrho_1 = \varrho_2 \oplus \mathbf{1}, \quad \varrho_2 \otimes \varrho_1 = \varrho_3 \oplus \varrho_1, \quad \text{etc.}$$

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

$$(2.36) \quad \varrho_m \otimes \varrho_n \simeq \varrho_{m+n} \oplus (\varrho_{m-1} \otimes \varrho_{n-1})$$

for $m \geq n \geq 1$.

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 as 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

¹⁴ In fact, any tensor product of finite-dimensional semisimple complex representations is semisimple, but this result of Chevalley is by no means easy to prove (see Theorem 7.1.11).

group $\mathrm{SL}_2(\mathbf{C})$ acts by simultaneous linear change of variable on the two sets of variables, i.e.,

$$(g \cdot f)(X_1, Y_1, X_2, Y_2) = f((X_1, Y_1)g, (X_2, Y_2)g)$$

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

$$\Delta \begin{cases} V_{m-1, n-1} & \longrightarrow & V_{m,n} \\ f & \mapsto & (X_1 Y_2 - X_2 Y_1) f \end{cases}$$

which is a G -homomorphism. If we view the factor $X_1 Y_2 - X_2 Y_1$ as 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 $d \geq 0$ with respect to some variable, say X_1 , and then $X_1 Y_2 f$ has degree $d + 1$ with respect to X_1 , while $X_2 Y_1 f$ remains of degree d , and therefore $X_1 Y_2 f \neq X_2 Y_1 f$.

Now we describe a stable complement to the image of Δ . To justify the solution a little bit, 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 at the linear map “evaluating f when both sets of variables are equal” suggested by the remark above, restricted to W . This map is given by

$$T \begin{cases} W & \longrightarrow & V_{m+n} \\ f & \mapsto & f(X, Y, X, Y) \end{cases}$$

(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 nonzero, 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 (Lemma 2.2.6) proves that T is surjective.

We now examine W more closely. Writing $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, we have

$$g \cdot e = (aX_1 + bY_1)^m (aX_2 + bY_2)^n = \sum_{0 \leq j \leq m+n} a^j b^{m+n-j} \varphi_j(X_1, Y_1, X_2, Y_2)$$

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 $\dim W \leq m + n + 1$. But since $\dim \varrho_{m+n} = m + n + 1$, we must have equality, and in particular T is an isomorphism.

Since $\dim V_{m-1,n-1} + \dim W = mn + m + n + 1 = \dim V_{m,n}$, it only remains to check that $V_{m-1,n-1} \oplus W = V_{m,n}$ to conclude that (2.36) 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 $f \in W \dots$ \square

In Corollary 5.6.2 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 centuries, 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 $m \geq 0$, 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 $d \geq 0$, 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 $S(V_m)_d^G$ 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*

$$\Delta(a_0, a_1, a_2) = a_1^2 - a_0 a_2$$

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 $S(V_3)$ the space of binary cubic forms, with elements

$$f = a_0 X^3 + 3a_1 X^2 Y + 3a_2 X Y^2 + a_3 Y^3 \in S(V_3),$$

one can prove that $S(V_3)^G \simeq \mathbf{C}[\Delta_3]$, where

$$\Delta_3 = a_0^2 a_3^2 - 6a_0 a_1 a_2 a_3 + 4a_0 a_2^3 - 3a_1^2 a_2^2 + 4a_1^3 a_3$$

(see, e.g., [59, 3.4.2 and 3.4.3]). We note that this reference computes the larger rings C_2 and C_3 of the *covariants* of quadratic and cubic binary forms, which have a double grading by integers $e \geq 0$ and $i \geq 0$, and that the invariants $S(V_2)^G$ and $S(V_3)^G$ are the $i = 0$ subrings of C_2 and C_3 , respectively.

The search for explicit descriptions of the invariant spaces $S(V_m)^G$ —and similar questions for other linear actions of groups such as $\mathrm{SL}_m(\mathbf{C})$ acting on homogeneous polynomials in more variables—was one of main topics of the classical theory of invariants, which was extremely popular during the 19th century (see, e.g., [59, Ch. 3] for a modern presentation). These questions are 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 $m \leq 10$. For $m = 9$, one needs 92 invariants to generate $S(V_m)^G$ as an algebra (see [10]); 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 $k(X)$ the k -vector space generated by basis vectors e_x indexed by the elements of the set X , and define

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

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 , which is called the *permutation representation* associated to the action of G on X . It has dimension $\dim \varrho = |X|$, by construction.

Example 2.6.4. (1) As the choice of notation suggests, the representation π_G of G on the space $k(G)$ spanned by G , defined in (2.3), 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

$$F = \{f : G \longrightarrow k \mid f(hg) = f(g) \text{ for all } h \in H\},$$

with the action of G given by the regular representation. This space has a basis given by the functions f_x which are the characteristic functions of the *left* cosets Hx . Moreover,

$$g \cdot f_x = f_{xg^{-1}}$$

(the left-hand side is non-zero at those y where $yg \in Hx$, i.e., $y \in Hxg^{-1}$), which means that mapping

$$f_x \mapsto e_{x^{-1}}$$

gives a linear isomorphism $F \longrightarrow k(X)$, which is now an intertwiner.

A feature of all permutation representations is that they are never irreducible if X is finite and $|X| \neq 1$: the element

$$\sum_{x \in X} e_x \in k(X)$$

is an invariant vector.

Exercise 2.6.5. (1) Let ρ be the permutation representation over a field k 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

$$\Phi \in \text{Hom}_G(\mathbf{1}, \rho)$$

constructed using Frobenius reciprocity.

(2) Let X be any finite set with an action of G , and let ρ be the associated permutation representation over k . Show that $\dim \rho^G$ is equal to the number of orbits of the action of G on X .

Exercise 2.6.6. Let $G = \mathfrak{S}_3$, and consider the permutation representation of dimension 3 associated to the natural action of G on $X = \{1, 2, 3\}$.

(1) Show that the subspace

$$E = \{v = (x, y, z) \in \mathbf{C}^3 \mid x + y + z = 0\}$$

is a stable complement of the line spanned by the invariant vector $(1, 1, 1) \in \mathbf{C}^3$ and that the representation ρ of G on E is irreducible.

(2) Show that $M = \mathbf{Z}^3 \subset \mathbf{C}^3$ and $N = E \cap M$ are stable lattices for their respective representations. Prove that the reduction of ρ modulo 3 (i.e., the representation of G over the field $\mathbf{Z}/3\mathbf{Z}$ induced from $N/3N$, as in the end of Section 2.4 and Exercise 2.4.3) is reducible and not semisimple.

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} describing 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

$$G \simeq \langle \underline{g} \mid \underline{r} \rangle$$

(the relations are complete in the sense that any relation between the generators is a product of *conjugates* of some of the given words.)

Then one can see that for a given field k and dimension $d \geq 1$, it is equivalent to give a d -dimensional (matrix) representation

$$G \longrightarrow \text{GL}_d(k)$$

or to give a family

$$(x_g)_{g \in \underline{g}}$$

of invertible matrices in $\text{GL}_d(k)$, 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 $\text{GL}_d(k)$.

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 relations to use, and since furthermore it might be very difficult to determine when representations built in this manner are isomorphic or not.

One very interesting well-known case of the use of generators and relations to define a representation is the construction of the *Weil representation* of $\mathrm{SL}_2(F)$, where F is a field which is either a finite field or \mathbf{R} . We refer to [41, Ch. XI] (for $F = \mathbf{R}$) or to [11, Prop. 4.1.3] (for finite fields) for full details, as well as to Exercise 4.6.21.

One can also make use of this approach to provide 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 $G \rightarrow \mathrm{GL}(E)$ 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 $\Phi : E \rightarrow F$ such that

$$x_g = \Phi^{-1} y_g \Phi$$

for all $g \in \underline{g}$.

Here is a slight variant that makes this even more 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

$$g_1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad g_2 = \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$$

in such a way that the only relations between the generators are

$$g_1^2 = 1, \quad g_2^3 = 1$$

(i.e., G is a free product of $\mathbf{Z}/2\mathbf{Z}$ and $\mathbf{Z}/3\mathbf{Z}$; see [26, II.A, II.B.28] for a proof and more information).

Hence it is equivalent to give a representation $\mathrm{PSL}_2(\mathbf{Z}) \rightarrow \mathrm{GL}(E)$ 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 and Baumslag; an example of a non-linear finitely generated group). *Let G be the group with two generators a, b subject to the*

relation

$$a^{-1}b^2a = b^3.$$

Then, whatever the field k , there exists no faithful linear representation

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

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 [45]) 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 (one says that G is *residually finite*; for a proof, see [46]).
- *The Identitätssatz of Magnus or Britton's Lemma* (see, e.g., [51, 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 faithful 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 $k \geq 1$. 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 relation is valid in any finite quotient.

Now Britton's Lemma [51, 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 .¹⁵ 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. \square

Remark 2.6.8. Concerning Malcev's Theorem, a good example to have 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 on 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 that some restrictions be imposed on the representations, namely, either we consider finite-dimensional representations or 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, let k be a field, and let*

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

be a k -representation of G .

¹⁵ In the language explained in Rotman's book [51], 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 $c \neq 1$.