

Lectures on the orbit method, by A. A. Kirillov, Graduate Studies in Mathematics, vol. 64, American Mathematical Society, Providence, RI, 2004, xx+408 pp., \$65.00, ISBN 0-8218-3530-0

1. INTRODUCTION WITHOUT FORMULAS

This book is about a wonderfully successful example of (if you will forgive some geometric language) circular reasoning. Here is the short version. Everywhere in mathematics, we find geometric objects M (like manifolds) that are too complicated for us to understand. One way to make progress is to introduce a vector space V of functions on M . The space V may be infinite-dimensional, but linear algebra is such a powerful tool that we can still say more about the function space V than about the original geometric space M . (With a liberal interpretation of “vector space of functions”, one can include things like the de Rham cohomology of M in this class of ideas.)

Often M comes equipped with a group G of symmetries, but G and M may be even less comprehensible together than separately. Nevertheless, G will act on our function space V by change of variables (giving linear transformations), and so we get a *representation* of G on V . Our original (and impossible) problem of understanding all actions of G on geometric spaces M is therefore at least related to the problem of understanding all representations of G . Because this is a problem about vector spaces, it sounds a bit less daunting.

The goal of the orbit method is to say something about all representations of a Lie group G . What it says is that an irreducible representation should correspond (roughly) to a symplectic manifold X with an action of G . So the circle is complete: we understand geometric objects with group actions in terms of representations, and we understand representations in terms of geometric objects with group actions.

That’s the end of the short version. Both the angel and the devil are in the details, so a longer version is called for. As an eternal optimist, I will begin with the angel. The geometric object at the end of the circle is a homogeneous symplectic manifold for G , with a little additional structure called in this book a *Poisson G -manifold*. (Kostant’s original terminology *Hamiltonian G -space* from [Ko] is also widely used.) Every homogeneous Poisson G -manifold is a covering of an orbit of the group G on the dual vector space \mathfrak{g}^* of its Lie algebra. The conclusion is that an irreducible representation of G (something interesting but difficult) should correspond to an orbit of G on \mathfrak{g}^* (something simple). This conclusion is the orbit method.

Here is an example. Suppose $G = GL(n, \mathbb{C})$, the group of invertible $n \times n$ matrices with complex entries. The dual of the Lie algebra of G may be identified with the vector space of *all* $n \times n$ matrices, where G acts by conjugation. The orbit method therefore says that irreducible representations of G should correspond to conjugacy classes of matrices: something that we teach undergraduates to parametrize, using Jordan canonical form.

So how is it that Jordan understood conjugacy classes of matrices in 1870 but that the last irreducible unitary representations of $GL(n, \mathbb{C})$ were found by Stein [St]

only in 1967? With this question the devil enters the story. The difficulty is with the word “should” in the phrase “a representation should correspond to an orbit”. In mathematics we expect such a correspondence to arise from a construction: a procedure to begin with an orbit and produce a representation, or perhaps to go in the other direction, or (preferably) both. For the orbit method, such procedures exist only for special groups or for special orbits in more general groups. Even when we have a list of orbits, we don’t immediately or in general have a list of “corresponding” representations.

A second problem is that representations do *not* correspond precisely to coadjoint orbits. In the case of $GL(n, \mathbb{C})$, the representations attached to orbits were found by Bargmann, Gelfand and Naimark, and others in the 1940s and 1950s, using methods developed by Mackey, Bruhat, and others ([B], [GN1], [GN2]). (This statement is historically misleading, since the idea of the orbit method appeared only in the 1960s.) These same authors found also the first examples of representations *not* attached to orbits: the “complementary series” representations of $SL(2)$. What Stein accomplished in his 1967 paper was to find the last of the representations not attached to orbits.

This seems to leave the orbit method as a kind of damaged treasure map, offering cryptic hints about where to find some (but certainly not all) of the representations we seek to understand. Why should we continue to consult it? First, parts of the map are perfect: for some groups (notably simply connected nilpotent Lie groups) there is a perfect bijection between orbits and irreducible unitary representations, given by explicit constructions and relating simple geometric properties to subtle representation-theoretic ones.

Second, the damaged parts of the map are sometimes the least interesting. A version of the Ramanujan conjecture says that any irreducible representation appearing in automorphic forms for $GL(n)$ (real or complex this time) must correspond to an orbit. (A kind of converse is implicit in the work of Langlands [La] in the 1960s: that a dense subset of representations of $GL(n)$ corresponding to orbits actually appears in automorphic forms.) This means that (conjecturally) the orbit method provides all the representations needed to study automorphic forms.

Third, the orbit method is the only map we have for representations of general Lie groups. In the case of reductive groups, the Langlands philosophy (specifically, Arthur’s refinement of the local Langlands conjecture) suggests that unitary representations should correspond to certain arithmetic objects. This alternative treasure map is (like the orbit method) extraordinarily powerful and useful. But it is more difficult to interpret than the orbit method, and it seems likely to lead in the end to a slightly smaller set of unitary representations.

I have so far addressed only the pragmatic question of whether the treasure map leads to treasure, neglecting entirely the question of why such a map should exist. Such neglect is natural in the real world, but for mathematicians it is strange. Kirillov’s book devotes more than two hundred pages to explaining representations and orbits separately, about a hundred and fifty to explaining how to read the map between them, and just over three pages to explaining why the orbit method ought to work. This is by no means an unreasonable distribution. If we actually knew why the orbit method worked, then a mathematics book on the subject might tell you only that, leaving the matter of examples and algorithms to the engineers. Since we don’t know, the only option is to write about examples and algorithms.

So why *does* the orbit method work? Kirillov offers two answers. The one that I find more convincing (probably because I understand it less) is due to Kostant [Ko] and to Souriau [So] and goes something like this. Symplectic geometry is a reasonable mathematical model for classical mechanics. The collection of all possible positions and momenta of all particles in a classical mechanical system (the phase space) is a symplectic manifold. Classical observables are functions on the phase space. An orbit—that is, a homogeneous Poisson G -manifold—may therefore be regarded as a classical mechanical system endowed with a group G of symmetries.

Hilbert space is a reasonable mathematical model for quantum mechanics. The collection of all (unnormalized) wave functions for a quantum mechanical system is a Hilbert space. Quantum observables are self-adjoint operators on that space. An irreducible representation may therefore be regarded as a quantum mechanical system endowed with a group G of symmetries.

Even though classical and quantum mechanics are different, they can sometimes be regarded as different descriptions of “the same” physical system. That is, for each classical mechanical system there should be a corresponding quantum mechanical system. This assertion is something less than physics or mathematics, yet it is not without content. At least sometimes, one can construct from a classical system the “corresponding” quantum system. In the presence of a group action, this construction—going from an orbit to an irreducible representation—is exactly what the orbit method says should exist.

The history of the orbit method is colorful and complicated, and I am not the one to sort it out properly. My own introduction to the subject came from Bert Kostant, and that will shape almost everything that I say. Because this is a review of a book by Kirillov, I will nevertheless try to use much of the language and viewpoint of the book. I have little doubt that I am going to make some attributions that are incomplete, incorrect, or even indefensible. For that I apologize. If you tell me about mistakes, I will try to correct them electronically and to make the same mistakes less often in the future.

Finally, I am very grateful for help from Tony Knapp with the relationship of Kirillov’s book to the work of Harish-Chandra.

2. INTRODUCTION WITH FORMULAS

To be more precise about the orbit method, we need to be more precise about the objects on each side of the correspondence. Representations are the complicated side in practice, but their definition is not too complicated, so I will begin there.

Definition 2.1. Suppose G is a topological group. A *unitary representation of G* is a pair (π, \mathcal{H}_π) with \mathcal{H}_π a Hilbert space and π a homomorphism from G to the group of unitary operators on \mathcal{H}_π . We assume that the map

$$G \times \mathcal{H}_\pi \rightarrow \mathcal{H}_\pi, \quad (g, v) \mapsto \pi(g)v$$

is continuous.

An *invariant subspace* for π is a closed subspace $W \subset \mathcal{H}_\pi$ with the property that $\pi(g)W \subset W$ for all $g \in G$. The representation is said to be *irreducible* if there are exactly two invariant subspaces (namely \mathcal{H}_π and 0).

It is helpful next to understand the special role of irreducible representations. If W is an invariant subspace of a unitary representation, then its orthogonal complement W^\perp is an invariant subspace as well, and we get a Hilbert space direct sum decomposition

$$\mathcal{H}_\pi = W \oplus W^\perp.$$

If \mathcal{H}_π is finite-dimensional, one can continue this process to get a decomposition of π as a Hilbert space direct sum of irreducible unitary representations. It is not very difficult to show that the collection of summands is uniquely determined (up to isomorphism) by the original representation π . That is, every finite-dimensional unitary representation can be written in just one way as a direct sum of irreducible unitary representations.

Trying to make a parallel construction in the case of infinite-dimensional representations leads to subtle limiting arguments. Many of the subtleties appear already in the theory of the Fourier transform on the real line. One way to think of that transform is as a decomposition of $L^2(\mathbb{R}, dx)$ as a “continuous direct sum” of the one-dimensional Hilbert spaces \mathbb{C}_ξ . The line \mathbb{C}_ξ has as orthonormal basis the function $f_\xi(x) = \exp(ix\xi)$; the difficulties arise because f_ξ does not belong to $L^2(\mathbb{R}, dx)$.

Theorem 2.2. *Suppose G is a separable locally compact group and π is a unitary representation of G . Then π is equivalent to a direct integral of irreducible unitary representations of G . If G is type I in the sense of von Neumann, then this direct integral decomposition is unique up to equivalence.*

I will not recall the definitions, except to say that $L^2(\mathbb{R}, dx)$ is the direct integral of the one-dimensional representations of \mathbb{R} on the spaces \mathbb{C}_ξ . (Left translation of f_ξ by $x \in \mathbb{R}$ multiplies f_ξ by the scalar $\exp(-ix\xi)$.) Many Lie groups are type I, including all real points of algebraic groups, but many solvable Lie groups are not type I. It is one of the unexpected successes of the orbit method (first found by Auslander and Kostant [AK] for simply connected solvable groups) that failure of G to be type I is sometimes manifest in bad topological behavior of the orbits of G on \mathfrak{g}^* .

Theorem 2.2 says that we can learn something about any unitary representation of G (including those on L^2 function spaces on subtle geometric spaces) by understanding irreducible unitary representations. I have not discussed in more detail what “understanding” might mean, and doing so would not bring us much closer to the orbit method. In the case of a Lie group G , one is often interested in group-invariant systems of differential equations on a manifold where G acts. Such systems can sometimes be analyzed using a direct integral decomposition from Theorem 2.2. In this case one needs in the end to understand the action of the universal enveloping algebra (or some special elements there) in each irreducible unitary representation of G .

We turn now to the orbit side of the orbit method. Although this is traditionally framed in the language of symplectic geometry, what arises directly in the orbit method is not the symplectic structure (a skew-symmetric bilinear form on tangent vectors) but the Poisson bracket (a Lie algebra structure on functions). I will therefore speak only about Poisson manifolds. These were introduced by Lie and studied in modern language by Hermann [He] and Lichnerowicz [Lz] (where the name originates). There is an introduction to the deep connections between symplectic and Poisson geometry in the beautiful article [We] of Weinstein.

Definition 2.3 (see [Ko, pages 176–177]). A *Poisson manifold* X is a smooth manifold endowed with a Poisson bracket

$$\{, \}: C^\infty(X) \times C^\infty(X) \rightarrow C^\infty(X),$$

subject to the following conditions.

- (1) The Poisson bracket makes $C^\infty(X)$ a Lie algebra. That is, it is bilinear, skew-symmetric, and satisfies the Jacobi identity

$$\{f, \{g, h\}\} = \{\{f, g\}, h\} + \{g, \{f, h\}\} \quad (f, g, h \in C^\infty(X)).$$

- (2) For each $f \in C^\infty(X)$, the endomorphism ξ_f of $C^\infty(X)$ defined by

$$\xi_f \cdot g = \{f, g\}$$

is a derivation:

$$\{f, gh\} = \{f, g\}h + g\{f, h\}.$$

We call ξ_f the *Hamiltonian vector field of f* . If we write $\text{Vect}(X)$ for the Lie algebra of smooth vector fields on X , then the Jacobi identity for the Poisson bracket says that the map

$$C^\infty(X) \rightarrow \text{Vect}(X), \quad f \mapsto \xi_f$$

is a Lie algebra homomorphism.

For the purposes of the orbit method, the most important example of a Poisson manifold is the vector space dual of a Lie algebra \mathfrak{g} . Each element $Z \in \mathfrak{g}$ defines a smooth (linear) function $f_Z \in C^\infty(\mathfrak{g}^*)$. The Poisson bracket is characterized by the requirement

$$(2.4) \quad \{f_Z, f_W\} = f_{[Z, W]} \quad (Z, W \in \mathfrak{g}).$$

Suppose that the Lie group G acts smoothly on the smooth manifold X . Differentiating the group action gives a Lie algebra homomorphism from the Lie algebra \mathfrak{g} of G to $\text{Vect}(X)$,

$$\mathfrak{g} \rightarrow \text{Vect}(X), \quad Z \mapsto \xi_Z.$$

Comparing this definition with Definition 2.3 suggests

Definition 2.5. A *Poisson G -manifold* is a Poisson manifold X endowed with

- (1) a smooth action of G , preserving the Poisson bracket, and
- (2) a G -equivariant Lie algebra homomorphism

$$\mathfrak{g} \rightarrow C^\infty(X), \quad Z \mapsto f_Z,$$

with the property that $\xi_Z = \xi_{f_Z}$ (Definition 2.3 and (2.4)).

For any finite-dimensional vector space V and any manifold X , a linear map from V to $C^\infty(X)$ is exactly the same thing as a smooth map from X to V^* . The mapping $Z \mapsto f_Z$ is therefore exactly the same thing as a smooth G -equivariant map of Poisson manifolds

$$\mu_X: X \rightarrow \mathfrak{g}^*.$$

The map μ_X is called the *moment map* for the Poisson G -manifold X .

Using the moment map, it is easy to deduce a classification for homogeneous Poisson G -manifolds. The result seems to be due independently to Kostant and to Kirillov; it appears as Lemma 1.4.6 in the book under review.

Theorem 2.6 ([Ki2, 15.2], and [Ko, Theorem 5.4.1]). *Suppose X is a homogeneous Poisson G -manifold. Then the moment map μ_X of Definition 2.5 is a G -equivariant covering map onto a single orbit of G on \mathfrak{g}^* .*

Conversely, if X is any homogeneous space for G and $\mu: X \rightarrow \mathfrak{g}^$ is any G -equivariant immersion, then X carries a unique structure of Poisson G -manifold with moment map μ .*

To repeat the main point: the orbit method suggests that (at least some) homogeneous Poisson G -manifolds should correspond to (at least some) irreducible unitary representations of G . The next question is which orbits correspond to which representations.

3. WHICH ORBITS COUNT?

As quantum mechanics was being developed, a fundamental notion was the integrality of certain numerical invariants. This integrality was imposed at first by hand as a way to recover some behavior seen in experiments. In later formulations the integrality appeared as discreteness of the spectrum of some Hilbert space operators, and so imposed by nature (or at least by the mathematical model).

In Kirillov's original application [Ki1] of the orbit method to a simply connected nilpotent Lie group, it happens that no discrete spectra arise in unitary representation theory. (This accident depends both on the simple connectivity and on the nilpotence of the group.) Consequently all orbits (of G on \mathfrak{g}^*) correspond to irreducible unitary representations.

As soon as one looks at more general Lie groups, operators with discrete spectrum appear, and unitary representations correspond only to orbits satisfying some sort of integrality condition. Making this integrality condition precise is a fundamental task, required even to read the orbit method as a treasure map. There are two schools of thought on how to do this, which I will label (to avoid the perils of engaging in history) "geometric" and "metaplectic". For nilpotent groups the two schools agree. For almost all other classes of groups they do not. I will explain roughly what the two methods say in a well-understood example, then say a few words about general definitions. Kirillov's book may be read as making the case for the geometric school; I will conclude this section with a few words from the metaplectic side.

Suppose that K is a compact connected Lie group with maximal torus T and Weyl group $W = W(K, T)$ (a finite group acting by automorphisms on T). Elementary structure theory for K provides a bijection

$$(3.1) \quad \text{orbits of } K \text{ on } \mathfrak{k}^* \leftrightarrow \text{orbits of } W \text{ on } \mathfrak{t}^*.$$

So the question is which W orbits on \mathfrak{t}^* should correspond to irreducible representations of K . (From here on I will make use of a choice of positive roots of T in $\mathfrak{k}_{\mathbb{C}}$, but it will mostly remain in the background.) Every irreducible unitary representation τ of K has a "highest weight", which is a one-dimensional character $\mu(\tau) \in \hat{T}$. The character $\mu(\tau)$ is determined by its differential, which is a linear map

$$(3.2)(a) \quad d\mu(\tau): \mathfrak{t} \rightarrow i\mathbb{R}.$$

Dividing by $2\pi i$, we get a linear functional

$$(3.2)(b) \quad \ell(\tau) = \frac{d\mu(\tau)}{2\pi i} \in \mathfrak{t}^*.$$

The linear functional $\ell(\tau)$ may reasonably still be called the highest weight of τ . Explicitly,

$$(3.2)(c) \quad \mu(\tau)(\exp Z) = \exp(2\pi i \ell(\tau)(Z)) \quad (Z \in \mathfrak{t}).$$

The geometric school attaches to τ the orbit of $\ell(\tau)$. Geometrically “integral orbits” correspond to those $\ell \in \mathfrak{t}^*$ such that $2\pi i \ell$ exponentiates to a character of T . This is a lattice in the vector space \mathfrak{t}^* .

On the other hand, every irreducible unitary representation τ of K is determined by its character $\Theta(\tau)$, which is a class function on K :

$$(3.3)(a) \quad \Theta(\tau)(k) = \text{tr } \tau(k).$$

A class function on K is determined by its restriction to T , which is a W -invariant function on T . Hermann Weyl gave a formula on T for the function $\Theta(\tau)$. Weyl’s formula involves a “Weyl denominator” Δ (which is independent of τ). The formula involves a linear functional $\lambda(\tau) \in \mathfrak{t}^*$ and looks like

$$(3.3)(b) \quad \Theta(\tau)(\exp(Z)) = \frac{\sum_{w \in W} \text{sgn}(w) \exp(2\pi i w \cdot \lambda(\tau)(Z))}{\Delta(\exp(Z))} \quad (Z \in \mathfrak{t}).$$

The metaplectic school attaches to τ the orbit of $\lambda(\tau)$. This differs from $\ell(\tau)$ by a certain translation ρ that is independent of τ :

$$\lambda(\tau) = \ell(\tau) + \rho.$$

Here ρ is half the sum of the positive roots of T in $\mathfrak{k}_{\mathbb{C}}$ (divided by $2\pi i$, like everything else). “Admissible orbits” in the metaplectic sense include all regular λ such that $\lambda - \rho$ exponentiates to a character of T . (There are additional singular λ which are also admissible in the metaplectic sense.)

This is a reasonable summary of the difference between the geometric and metaplectic approaches: whether a representation of a compact group should correspond to its highest weight or to an exponent in the Weyl character formula.

Here is how this idea is made precise for general Lie groups. In each case the idea is that a representation should correspond to an orbit together with some additional structure and that this additional structure can exist only when the orbit has some integrality property. Here is the geometric version as it appears in Kirillov’s book. (One can find a beautiful justification for the term “geometric” in [Ko, Corollary 1 to Theorem 5.7.1].)

Definition 3.4 (see the book under review, page 123). Suppose G is a Lie group and $F \in \mathfrak{g}^*$. Write $G^F \subset G$ for the stabilizer of F in the action on \mathfrak{g}^* , so that the orbit of F may be identified with the homogeneous space G/G^F . Write $\mathfrak{g}^F \subset \mathfrak{g}$ for the Lie algebra of the stabilizer of F . It turns out that the restricted linear functional

$$F: \mathfrak{g}^F \rightarrow \mathbb{R}$$

is automatically a one-dimensional Lie algebra representation. This Lie algebra representation may or may not exponentiate to the identity component of the group G^F . A *rigged momentum* at F is by definition an irreducible unitary group representation (ϕ, \mathcal{H}_ϕ) of G^F , with the property that

$$d\phi(Z) = 2\pi i F(Z) \cdot I_\phi \quad (Z \in \mathfrak{g}^F).$$

Here I_ϕ is the identity operator on \mathcal{H}_ϕ . The existence of rigged momenta at F is an integrality constraint on F , and their uniqueness is controlled by the disconnectedness of G^F . A *rigged coadjoint orbit* is an orbit of G on rigged momenta.

The geometric version of the orbit method is that each rigged coadjoint orbit should correspond to an irreducible unitary representation of G . This is exactly correct for simply connected nilpotent groups (where all the groups G^F are connected and simply connected, so that a rigged coadjoint orbit is the same thing as a coadjoint orbit). It is true for simply connected type I solvable groups, by the results of Auslander and Kostant in [AK]. It is true for compact Lie groups (even disconnected ones) by the Cartan-Weyl highest weight theory.

Here is the metaplectic version. I will not try to trace its history; in this form it is due to Duflo.

Definition 3.5 (see for example [GV], Definition 6.2). Suppose G is a Lie group and $F \in \mathfrak{g}^*$. The Poisson structure on the orbit of F defines a symplectic structure ω_F on the vector space $T_F = \mathfrak{g}/\mathfrak{g}^F$ (the tangent space at F to the orbit). Write $Sp(T_F)$ for the group of linear transformations of T_F preserving ω_F . The adjoint action defines a map

$$G^F \rightarrow Sp(T_F).$$

The symplectic group $Sp(T_F)$ has a double cover (the *metaplectic group*)

$$1 \rightarrow \{1, \epsilon\} \rightarrow Mp(T_F) \rightarrow Sp(T_F) \rightarrow 1,$$

arising from the interpretation of $Sp(T_F)$ as acting on the canonical commutation relations (see [GV]). Pulling this covering back to G^F gives

$$1 \rightarrow \{1, \epsilon\} \rightarrow \tilde{G}^F \rightarrow G^F \rightarrow 1.$$

An *admissible orbit datum* at F is by definition an irreducible unitary group representation (ϕ, \mathcal{H}_ϕ) of \tilde{G}^F , with the property that

$$d\phi(Z) = 2\pi i F(Z) \cdot I_\phi \quad (Z \in \mathfrak{g}^F), \quad \text{and} \quad \phi(\epsilon) = -I_\phi.$$

The existence of admissible orbit data at F is a kind of “half-integrality” constraint on F , and their uniqueness is controlled by the disconnectedness of G^F .

The metaplectic version of the orbit method is that each admissible orbit datum should correspond to an irreducible unitary representation of G . This is precisely correct for simply connected nilpotent groups, where again there is a unique admissible orbit datum on each orbit. It is true for compact Lie groups, even disconnected ones, if one substitutes “irreducible or zero” for “irreducible”, but several different admissible orbit data may correspond to the same representation. (For example, the trivial representation corresponds both to the singular orbit $\{0\}$ and to the regular orbit of ρ .) Duflo’s definition is designed to work better than rigged momenta for solvable groups that are not simply connected, but I do not know whether the “admissible data” version of [AK] holds in that generality.

For me there is compelling evidence for the metaplectic approach in the case of an elliptic coadjoint orbit for a reductive group G (see for example [V, Lecture 3]). For such an orbit, there is a natural construction of a (possibly reducible) unitary representation attached to each admissible orbit datum. Harish-Chandra’s description of the discrete series may be interpreted to say that the discrete series of G is parametrized by admissible orbit data on regular elliptic orbits.

Suppose we now consider the restriction of a discrete series representation to a maximal compact subgroup K of G . On the level of representations this restriction is computed by the Blattner conjecture, proved by Hecht and Schmid. On the level of orbits there is a geometric prediction in Kirillov's Rule 3 (from the introduction to his book). In the language of admissible orbit data the representations and the geometry correspond perfectly; indeed, Duflo, Heckman, and Vergne gave in [DHV] a *proof* of Blattner's conjecture based on the orbit geometry and Kirillov's "universal character formula".

The notions of admissible and rigged orbit are quite different for elliptic orbits. I believe it is a straightforward exercise to make examples (based on the Blattner formula) showing that Kirillov's prediction of restriction to K —Rule 3 of the introduction to the book under review—cannot be realized for discrete series using rigged elliptic orbits.

4. KIRILLOV'S BOOK

Kirillov's book seeks to introduce young mathematicians to the orbit method. It is aimed at beginning graduate students in the U.S. Two hundred pages of appendices recall or introduce the notions of smooth manifold, Lie group, category, and much more. There are hundreds of examples, well chosen to illustrate the ideas. The orbit method itself is introduced slowly and carefully: first on nilpotent groups (where it works perfectly), then on solvable groups (where the difficulties are carefully explained in the simplest possible examples), and then on compact groups. Readers already familiar with (say) representation theory for compact groups will have an easier time, but those who are not will find all the explanations they need. The philosophical underpinnings of this text are perfect: this is the way everyone should try to write graduate texts.

But... I have a serious problem with the way the text was refereed and edited. A fundamental characteristic of mathematics is that it can make statements that are exactly and completely correct. This is a wonderful tool for learning a subject: if you find the smallest formal inconsistency between two things that the text seems to say, then you know that you have missed something. You can read theorems about differential operators in a functional analysis text and know that they apply to the differential operators you use for representation theory.

You can do all of that if the book is carefully put together. This book is not. The results about representations and coadjoint orbits are mostly true, but the more general mathematical statements from which they are derived are often false. One example is on page 225, where it is asserted that $H_0(X, \mathbb{Z}) = \pi_0(X)$ for any topological space X . For anyone who has taken an algebraic topology course, this is easy to fix, but the formula is given exactly for readers who have *not* taken such a course. A more serious example appears in Proposition 2 on page 229. This asserts that if G acts smoothly on a manifold M and all orbits are of the same dimension, then M/G is a manifold. This is false: a dull example is the action of $\mathbb{Z}/2\mathbb{Z}$ on the circle by turning it over, but there are more subtle ones (with G connected, for example).

Since I support the goal of attracting young mathematicians to the orbit method and this book could contribute to that, I have made a list of corrections (along with a few philosophical objections). It is available at

<http://www-math.mit.edu/~dav/CORR.pdf>.

I welcome additions and corrections and will try to update the list accordingly.

REFERENCES

- [AK] L. Auslander and B. Kostant, *Polarization and unitary representations of solvable Lie groups*, Invent. Math. **14** (1971), 255–354. MR0293012 (45:2092)
- [B] V. Bargmann, *Irreducible unitary representations of the Lorentz group*, Ann. of Math. **48** (1947), 568–640. MR0021942 (9:133a)
- [DHV] M. Duflo, G. Heckman, and M. Vergne, *Projection d’orbites, formule de Kirillov et formule de Blattner*, Harmonic analysis on Lie groups and symmetric spaces (Kleebach, 1983), Mém. Soc. Math. France (N.S.), No. 15, 1984, pp. 65–128. MR0789081 (86m:22017)
- [GN1] I. M. Gelfand and M. A. Naimark, *Unitary representations of the Lorentz group*, Izv. Akad. Nauk S.S.S.R. **11** (1947), 411–504. MR0024440 (9:495a)
- [GN2] I. M. Gelfand and M. A. Naimark, *Unitary Representations of the Classical Groups*, Trudy Mat. Inst. Steklov, vol. 36, Moscow-Leningrad, 1950; German transl.: Akademie-Verlag, Berlin, 1957. MR0046370 (13:722f)
- [GV] W. Graham and D. Vogan, *Geometric quantization for nilpotent coadjoint orbits*, Geometry and Representation Theory of Real and p -adic Groups (J. Tirao, D. Vogan, and J. Wolf, eds.), Birkhäuser, Boston-Basel-Berlin, 1998, pp. 69–137. MR1486137 (2000i:22024)
- [He] R. Hermann, *Toda lattices, cosymplectic manifolds, Bäcklund transformations and kinks. Part A.*, Interdisciplinary Mathematics, Vol. XV, Math. Sci. Press, Brookline, MA, 1977. MR0478194 (57:17682)
- [Ki1] A. Kirillov, *Unitary representations of nilpotent Lie groups*, Uspehi Mat. Nauk. **17** (1962), 57–110. MR0142001 (25:5396)
- [Ki2] A. Kirillov, *Elements of the Theory of Representations*, English transl. by E. Hewitt, Springer-Verlag, Berlin-Heidelberg-New York, 1976. MR0412321 (54:447)
- [Ko] B. Kostant, *Quantization and unitary representations*, Lectures in Modern Analysis and Applications (C. Taam, ed.), Lecture Notes in Mathematics, vol. 170, Springer-Verlag, Berlin-Heidelberg-New York, 1970, pp. 87–208. MR0294568 (45:3638)
- [La] R. P. Langlands, *On the functional equations satisfied by Eisenstein series*, Lecture Notes in Mathematics, vol. 544, Springer-Verlag, Berlin-New York, 1976. MR0579181 (58:28319)
- [Lz] A. Lichnerowicz, *Les variétés de Poisson et leurs algèbres de Lie associées*, J. Differential Geometry **12** (1977), 253–300. MR0501133 (58:18565)
- [So] J.-M. Souriau, *Structure des systèmes dynamiques*, Maîtrises de mathématiques, Dunod, Paris, 1970; English transl. by C. H. Cushman-de Vries in Structure of Dynamical Systems: A Symplectic View of Physics, Progress in Mathematics, vol. 149, Birkhäuser Boston, Boston, MA, 1997. MR0260238 (41:4866)
- [St] E. M. Stein, *Analysis in matrix spaces and some new representations of $SL(n, \mathbb{C})$* , Ann. of Math. (2) **86** (1967), 461–490. MR0219670 (36:2749)
- [V] D. Vogan, *The method of coadjoint orbits for real reductive groups*, Representation Theory of Lie Groups, IAS/Park City Mathematics Series, vol. 8, American Mathematical Society, Providence, RI, 2000. MR1737729 (2001k:22027)
- [We] A. Weinstein, *The local structure of Poisson manifolds*, J. Differential Geom. **18** (1983), 523–557. MR0723816 (86i:58059)

DAVID A. VOGAN, JR.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

E-mail address: dav@math.mit.edu