The method of coadjoint orbits for real reductive groups

David A. Vogan, Jr.

IAS/Park City Mathematics Series
Volume 6, 1998

The method of coadjoint orbits for real reductive groups

David A. Vogan, Jr.

Introduction

The method of coadjoint orbits suggests that irreducible unitary representations of a Lie group G are something like quantum mechanical systems, and that their classical analogues are (roughly speaking) symplectic homogeneous spaces. (Actually one imposes on the symplectic homogeneous space two additional structures: that of a Hamiltonian G-space, and an admissible orbit datum. We will ignore these for the purposes of the introduction.) The problem of quantization in mathematical physics is to attach a quantum mechanical model to a classical physical system. The problem of quantization in representation theory is to attach a unitary group representation to a symplectic homogeneous space.

The miraculous aspect of the orbit method is twofold. First, the relevant symplectic homogeneous spaces can easily be classified: they are the covering spaces of the orbits of G on the dual of its own Lie algebra. (This is the source of the terminology "coadjoint orbits.") Second, the method works: it has successfully suggested where to look for large families of very different unitary representations of very different Lie groups.

Both of these claims of miraculousness require some justification. For the possibility of classifying symplectic homogeneous spaces, one can think of the superficially similar problem of classifying Riemannian homogeneous spaces. Even for as simple a group as $SO(3)$, the answer to this problem is quite complicated (see Exercise 1). For a group like $GL(n,\mathbb{R})$, I do not know any reasonable description of the Riemannian homogeneous spaces. Yet we will see that the symplectic homogeneous spaces for any semisimple Lie group can be parametrized quite precisely in fairly elementary terms.

That the effectiveness of the method is miraculous—that is, that we should be surprised at a close connection between irreducible unitary representations and symplectic homogeneous spaces—requires even more justification. Mathematics

[1]Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
E-mail address: dav@math.mit.edu

is full of bijections between sets that at first glance appear unrelated (like commutative Banach algebras and compact topological spaces, for example; or vector spaces and cardinal numbers). We learn to expect that there is a clever idea that builds a bridge, and that once we have mastered the bridge, we can move back and forth freely and easily. In the case of the orbit method, there seems to be no such bridge. Relatively simple groups like $SL(2,\mathbb{R})$ have irreducible unitary representations (the complementary series) that do not correspond to any symplectic homogeneous space. Conversely, Torasso found in [**Tor**] that the double cover of $SL(3,\mathbb{R})$ has a symplectic homogeneous space corresponding to no unitary representation. (The space in Torasso's counterexample actually carries the more refined structure of Hamiltonian G-space and admissible orbit datum mentioned earlier.)

If there is no big theorem at the end of the orbit method's rainbow, why should we be chasing it? The point is that unitary representations remain extraordinarily mysterious. The orbit method provides some light in a very dark room. Many large families of coadjoint orbits do correspond in comprehensible ways to unitary representations, and provide a clear geometric picture of these representations. Many ideas about orbits (like the Jordan decomposition in the case of reductive groups) suggest corresponding properties of unitary representations; and these properties can sometimes be proved. Best of all, coadjoint orbits can tell us where to look for unitary representations that we haven't yet thought of.

The first goal of these notes is to describe what is known about the quantization problem for reductive Lie groups, and the extent to which its resolution can lead to all unitary irreducible representations. We will find that the quantization problem comes down to a finite set of coadjoint orbits for each reductive group: the nilpotent orbits.

The problem of quantizing nilpotent orbits has a long and colorful history. Two themes recur: quantization of very particular classes of orbits, and incomplete schemes for quantization of general orbits. I will mention one favorite example of each: Torasso's paper [**Tor2**] on quantization of minimal orbits, and a joint paper [**GV**] with Graham suggesting how to guess at the representations attached to general orbits.

The second goal of the notes is to describe a new method for studying the quantization of nilpotent coadjoint orbits in terms of restriction to a maximal compact subgroup. This falls into the second "recurring theme," and therefore the results will be very incomplete. What I intend to show is that the problem of quantizing a particular nilpotent orbit can often be reduced to a finite amount of linear algebra; and that there may be hope for understanding this linear algebra in general.

I am grateful to the graduate students at the 1998 PCMI for their interest and for their critical comments on these notes. Most of all I would like to thank Monica Nevins, who served as teaching assistant for the summer course and read the notes with great care. Her mathematical advice is responsible for many pockets of clarity in the exposition. (After her last corrections, I sprinkled a number of typographical and mathematical absurdities through the text, in order to restore it to my usual standards. For these she bears no responsibility.)

LECTURE 1
Some ideas from mathematical physics

The method of coadjoint orbits has its origins in mathematical physics. As explained for example in [**Arn**], a classical mechanical system can often be modelled by a symplectic manifold X, called the phase space. A point of the manifold records something like the positions and momenta of all the particles in the system. The evolution of the system in time defines a path γ in X. Newton's laws for the evolution of the system say that this path is an integral curve for a Hamiltonian vector field ξ_E on X:

$$d\gamma/dt = \xi_E(\gamma(t)). \tag{1.1)(a}$$

The function E on X corresponds to the energy of the mechanical system. Other physical observables correspond to other functions f on X. Along an integral curve γ for ξ_E, the observable f evolves according to the differential equation

$$d(f \circ \gamma)/dt = \{E, f\} \circ \gamma; \tag{1.1)(b}$$

here $\{E, f\}$ is the Poisson bracket of E and f. In particular, the observable f is *conserved* (that is, constant in time) if and only if it Poisson commutes with the energy.

A quantum mechanical system, on the other hand, is typically modelled by a Hilbert space $\mathcal{H}$. Each state of the system corresponds to a line in the Hilbert space. The evolution of the system in time therefore corresponds to a path in the projective space $\mathbb{P}\mathcal{H}$. The physical laws for the system are encoded by a one-parameter group of unitary operators, whose orbits on $\mathbb{P}\mathcal{H}$ are the possible time histories of the system. The generator of this one-parameter group is a skew-adjoint operator A, again corresponding to the energy of the system. (We will be careless about questions of unbounded operators, domains, and so on; these questions seem not to affect the philosophical ideas we want to emphasize.) If $v(t)$ is a vector in $\mathcal{H}$ representing the state of the system at time t, then the analogue of (1.1)(a) is the Schrödinger equation

$$dv/dt = Av(t). \tag{1.2)(a}$$

(This will look more like a classical Schrödinger equation in the setting of (1.3)(b) below.) In general a physical observable corresponds to an operator B on $\mathcal{H}$. If the system is in a state corresponding to a line L in $\mathbb{P}\mathcal{H}$, then the outcome of the observation corresponding to B cannot be predicted exactly. There is a probability distribution of possible outcomes, with expectation $\langle Bv, v\rangle$; here v is any unit vector in L. Along an orbit of our one-parameter group e^{tA} of unitary operators, this expectation becomes

$$\langle e^{-tA}Be^{tA}v, v\rangle. \tag{1.2(b)}$$

Consequently B is conserved exactly when it commutes with the one-parameter group. A more precise analogue of (1.1)(b) is the differential equation

$$d/dt(\langle Bv, v\rangle) = \langle (BA - AB)v(t), v(t)\rangle. \tag{1.2(c)}$$

Again the conserved operators are those commuting with the energy.

These two kinds of mathematical model are fundamentally different. Asking for any close correspondence between them seems to fly in the face of all the painfully learned lessons of physics in this century. Nevertheless, such correspondences sometimes exist. That is, certain classical mechanical systems correspond formally to quantum mechanical ones. Because such correspondences are at the heart of the orbit philosophy, we need to understand them in some detail.

A simple case (arising for example from the motion of point particles) is a system attached to a smooth Riemannian manifold M of "states" (typically the possible positions of the particles in the system). A tangent vector to M at a state m specifies the velocities of the particles; the Riemannian metric gives kinetic energy. The phase space X of such a mechanical system may be identified with the cotangent bundle $T^*(M)$, endowed with its natural symplectic structure. The energy function E on $T^*(M)$ may be the sum of a kinetic energy term (the Riemannian length in the cotangent direction) and a potential energy term depending only on M (that is, on the positions of the particles). Explicitly,

$$E(m, \xi) = \langle \xi, \xi\rangle + V(m) \qquad (m \in M, \xi \in T^*_m(M)). \tag{1.3(a)}$$

There is a simple and natural family of observables in this system. If q is any smooth function on M, then q defines a function on the phase space $T^*(M)$; the value of this function depends only on the positions of the particles in the system, and not on their velocities. One sees easily that any two such functions on $T^*(M)$ must Poisson commute with each other:

$$\{q_1, q_2\} = 0 \qquad (q_i \in C^\infty(M)). \tag{1.3(b)}$$

Such a system has a natural quantization. The Hilbert space $\mathcal{H}$ is the complex Hilbert space $L^2(M)$. If v is a unit vector in this Hilbert space, then $|v|^2$ is a non-negative function on M of integral 1; that is, a probability density on M. One thinks of this density as describing the probability of finding the particles of the system in a particular position. The quantum-mechanical energy operator is the Laplace-Beltrami operator L on M, plus the potential V. More precisely,

$$A = i\hbar L + \frac{i}{\hbar}V. \tag{1.3(c)}$$

(The factor $\hbar$ is Planck's constant divided by 2π.) There is a natural family of quantum observables corresponding to the classical position observables. If q is any smooth function on M, then multiplication by q defines an operator A_q on $L^2(M)$. Any two of these operators commute with each other:

$$[A_{q_1}, A_{q_2}] = 0 \qquad (q_i \in C^\infty(M)). \tag{1.3(d)}$$

(A little thought shows that the identification of these observables as "positions" is more or less equivalent to our earlier assertion that the probability density $|v|^2$ describes the distribution across M of the quantum state.)

Physically the idea of quantization is to replace a classical mechanical model by a quantum mechanical model of the "same" system. This idea is a little difficult to pin down, since part of the point of quantum mechanics is that there never was a physically real classical system in the first place. Mathematically the idea is to pass from a symplectic manifold X and an interesting collection of real-valued functions $\{f_i \mid i \in I\}$ on X (the classical observables), to a Hilbert space $\mathcal{H}$ and an interesting collection of skew-adjoint operators $\{A_i \mid i \in I\}$ on $\mathcal{H}$ (the quantum observables). The shape of the equations (1.1)(b) and (1.2)(c) for the time evolution of observables suggests that we might ask for quantization to carry Poisson bracket of functions to commutator of operators. That is,

$$\{f_i, f_j\} = f_k \Rightarrow [A_i, A_j] = A_k. \tag{1.4}$$

The requirement (1.4) turns out to be too stringent to impose on all observables. Roughly speaking, the reason is that functions under Poisson bracket are more nearly commutative than operators under commutator. To understand this statement, it is helpful to look at an example. Let us take for X the $2n$-dimensional vector space $\mathbb{R}^{2n}$, endowed with the standard symplectic structure

$$\omega((x,\lambda),(y,\xi)) = \langle x, \xi\rangle - \langle y, \lambda\rangle \qquad (x,y,\xi,\lambda \in \mathbb{R}^n). \tag{1.5(a)}$$

An interesting family of observables consists of the affine functions on X. As a basis of these, we can take the constant function 1 together with the linear functions

$$p_i(x,\lambda) = x_i, \qquad q_j(x,\lambda) = \lambda_j. \tag{1.5(b)}$$

The constant function Poisson commutes with all functions, and the linear functions satisfy

$$\{p_i, p_{i'}\} = 0, \qquad \{q_j, q_{j'}\} = 0, \qquad \{p_i, q_j\} = \delta_{ij}\cdot 1. \tag{1.5(c)}$$

The quantization problem as formulated at (1.4) asks us to realize these same relations as commutators of self-adjoint operators on a Hilbert space. There is a very natural way to do that. We define

$$\mathcal{H} = L^2(\mathbb{R}^n), \tag{1.6(a)}$$

thinking of this as the $\mathbb{R}^n$ in the last n coordinates of X. The functions q_j are then naturally identified with the coordinate functions on $\mathbb{R}^n$. Multiplication by such a coordinate function is a self-adjoint operator on L^2, so it is reasonable to define

$$Q_j = \text{multiplication by } \sqrt{-1}q_j, \tag{1.6(b)}$$

a skew-adjoint operator on $L^2(\mathbb{R}^n)$. Similarly, we put

$$P_i = \partial/\partial q_i, \tag{1.6(c)}$$

a skew-adjoint operator on L^2. These operators satisfy commutation relations

$$[P_i, P_{i'}] = 0, \qquad [Q_j, Q_{j'}] = 0, \qquad [P_i, Q_j] = \delta_{ij}\cdot\sqrt{-1}. \tag{1.6(d)}$$

If we make the constant function 1 correspond to the scalar quantum operator $\sqrt{-1}$, then the requirement of (1.4) is satisfied.

The problem with (1.4) becomes apparent when we try to quantize more functions. A natural extension of the class of affine functions on X is the class of polynomial functions. This is just the polynomial algebra with generators p_i and q_j:

$$\mathrm{Poly}(X) = \mathbb{R}[p_1, \dots, p_n, q_1, \dots, q_n]. \tag{1.7)(a}$$

Similarly, the span of the operators P_i, Q_j, and the imaginary scalars extends naturally to the skew-adjoint polynomial coefficient differential operators:

$$\mathrm{Diff}(\mathbb{R}^n) = \{D \in \mathbb{C}[\partial/\partial q_1, \dots, \partial/\partial q_n, q_1, \dots, q_n] \mid D^* = -D\}. \tag{1.7)(b}$$

There are many ways to arrange vector space isomorphisms between $\mathrm{Poly}(X)$ and $\mathrm{Diff}(\mathbb{R}^n)$, but none that respect the two kinds of commutator (see Exercise 3).

One lesson that might be drawn from the details of Exercise 3 is that finite-dimensional Lie algebras of classical observables have a better chance to quantize nicely. This restriction does not interfere with our particular interest in representation theory of finite-dimensional Lie groups. As a first approximation to a more mathematical notion of quantization, we may think of a symplectic manifold X endowed with a finite-dimensional Lie algebra $\mathfrak{g}$ of smooth functions. A quantization would then be a Hilbert space endowed with a finite-dimensional Lie algebra (also isomorphic to $\mathfrak{g}$) of skew-adjoint operators. This is close to where we want to be; but we need to face some of the problems about unbounded operators that we have ignored until now. On X, the Lie algebra of functions defines a Lie algebra of (Hamiltonian) vector fields; the question there is whether these vector fields may be integrated to a group action on X. On the Hilbert space, the question is whether the family of unbounded skew-adjoint operators can be integrated to a group representation. Both questions are subtle and difficult, but to some extent irrelevant: we are concerned finally only with cases when the integration is possible. Here is the classical setting.

Definition 1.8. Suppose X is a symplectic manifold, and $f \in C^\infty(X)$ is a smooth function. The *Hamiltonian vector field* ξ_f *of* f is defined in terms of the Poisson bracket by

$$\xi_f(g) = \{f, g\}.$$

Suppose G is a Lie group endowed with a smooth action on X by symplectomorphisms. We say that X is a *Hamiltonian G-space* if there is a linear map

$$\tilde{\mu}: \mathfrak{g} \to C^\infty(X)$$

with the following properties. First, $\tilde{\mu}$ intertwines the adjoint action of G on $\mathfrak{g}$ with its action on $C^\infty(X)$. Second, for each $Y \in \mathfrak{g}$, the vector field by which Y acts on X is $\xi_{\tilde{\mu}(Y)}$. Third, $\tilde{\mu}$ is a Lie algebra homomorphism.

This definition may sound circular. In most developments of symplectic manifolds the Hamiltonian vector field ξ_f is defined first (by using the symplectic form to identify the 1-form df with a vector field). Then the formula in Definition 1.8 is used as the definition of the Poisson bracket. But for the theory of quantization the Poisson bracket is more fundamental than the symplectic form, so it is not unreasonable to define other ideas in terms of the Poisson bracket. One consequence is that Definition 1.8 can be formulated in the category of Poisson manifolds, or even of (possibly singular) Poisson algebraic varieties. The definition is due essentially to Kirillov and to Kostant independendently (see [**Kir, 15.2**] and [**Kos, 5.4.1**]).

The natural quantum analogue of a Hamiltonian G-space is simply a unitary representation of G.

Definition 1.9. Suppose G is a Lie group. A *unitary representation of* G is a pair $(\pi, \mathcal{H})$ with $\mathcal{H}$ a Hilbert space, and

$$\pi\colon G \to U(\mathcal{H})$$

a homomorphism from G to the group of unitary operators on H. That is, we require a continuous action map

$$G \times \mathcal{H} \to \mathcal{H}$$

that preserves lengths in the Hilbert space.

To make a stronger connection with the formalism of Definition 1.8, we should introduce the Lie algebra representation $d\pi$ of $\mathfrak{g}$; the operators of $d\pi$ are defined on a dense subspace $\mathcal{H}^\infty$ of the Hilbert space, and they are formally skew-adjoint. But making correct statements about the relationship between these operators and π is difficult and unnecessary; so we have simply dropped them. The reader may like to consider whether Definition 1.8 might similarly be modified to avoid any mention of the Lie algebra.

The idea that we want from mathematical physics is that there should be a notion of "quantization" passing from Definition 1.8 to Definition 1.9: that is, from Hamiltonian G-spaces to unitary representations. We will sharpen this idea in various ways as we go along.

One sharpening is available immediately. Definitions 1.8 and 1.9 both admit simplest cases from which more general ones might be built. In the case of Definition 1.8, these are the *homogeneous* Hamiltonian G-spaces; that is, those for which the action on X is transitive. In the case of Definition 1.9, it is the slightly more subtle notion of *irreducibility*: that $\mathcal{H}$ should admit exactly two closed subspaces invariant under the action of G. We would like quantization to carry homogeneous Hamiltonian G-spaces to irreducible unitary representations.

LECTURE 2

The Jordan decomposition and three kinds of quantization

As explained in the introduction, the first miraculous aspect of the orbit method is that the homogeneous Hamiltonian G-spaces are easy to classify. The result is due to Kirillov and to Kostant. Here is the basic construction.

Lemma 2.1 **([Kir], 15.2, [Kos], Theorem 5.4.1).** *Suppose G is a Lie group with Lie algebra $\mathfrak{g}$, and $X \subset \mathfrak{g}^*$ is an orbit of the coadjoint action of G. For $\lambda \in X$, write G^λ for the stabilizer of λ; then $X \simeq G/G^\lambda$.*

1. *The tangent space to X at λ is*

$$T_\lambda(X) \simeq \mathfrak{g}/\mathfrak{g}^\lambda.$$

2. *The skew-symmetric bilinear form*

$$\omega_\lambda(u,v) = \lambda([u,v]) \qquad (u, v \in \mathfrak{g})$$

on $\mathfrak{g}$ has radical exactly $\mathfrak{g}^\lambda$, and so defines a symplectic form on $T_\lambda(X)$.

3. *The forms ω_λ make X into a symplectic manifold.*
4. *For $u \in \mathfrak{g}$, define $\tilde{\mu}(u) \in C^\infty(X)$ to be the restriction to X of the linear functional u on $\mathfrak{g}^*$. Then the mapping $\tilde{\mu}$ makes X a Hamiltonian G-space (Definition 1.8).*

Lemma 2.2. *Suppose that X is a Hamiltonian G-space (Definition 1.8). Suppose $\widetilde{X}$ is another smooth manifold with a G action, and that $\pi: \widetilde{X} \to X$ is a G-equivariant local diffeomorphism. Then $\widetilde{X}$ inherits from X a natural structure of Hamiltonian G-space.*

This is entirely elementary, and we leave it to the reader.

Theorem 2.3 **([Kir], 15.2, [Kos], Theorem 5.4.1).** *The homogeneous Hamiltonian G-spaces for a Lie group G are the covering spaces of coadjoint orbits. More precisely, suppose X is such a space, with moment map $\tilde{\mu}: \mathfrak{g} \to C^\infty(X)$. Reinterpret $\tilde{\mu}$ as a map*

$$\mu: X \to \mathfrak{g}^*, \qquad \mu(x)(u) = \tilde{\mu}(u)(x) \quad (x \in X, u \in \mathfrak{g}).$$

Then μ is a G-equivariant local diffeomorphism onto a coadjoint orbit, and the Hamiltonian G-space structure on X is pulled back from that on the orbit (Lemma 2.1) by the map μ (Lemma 2.2).

To get a typical homogeneous Hamiltonian G-space, we must therefore begin with an element $\lambda \in \mathfrak{g}^*$, and its stabilizer G^λ. A covering of the orbit is a homogeneous space G/G_1^λ, with

$$G_0^\lambda \subset G_1^\lambda \subset G^\lambda.$$

The subgroup G_1^λ is determined by the space up to conjugation in G^λ. These observations are summarized in

Corollary 2.4. *The homogeneous Hamiltonian G-spaces covering the coadjoint orbit $G \cdot \lambda$ are in one-to-one correspondence with conjugacy classes of subgroups of the discrete group G^λ/G_0^λ of connected components of G^λ.*

In order to study quantization for a Lie group G, we must therefore first understand the orbits of G on $\mathfrak{g}^*$.

We are particularly interested in reductive Lie groups. Once it was traditional to look at connected semisimple Lie groups with finite center, but now it is clear that inductive arguments require a more flexible class: some abelian factors should be allowed, and some disconnectedness. Exactly which groups should be allowed is a surprisingly delicate question. The many pleasant properties of connected semisimple groups disappear not all at once but one or two at a time as the hypotheses are weakened. Harish-Chandra introduced a class of reductive groups that works well for many purposes. On the other hand his class excludes even some disconnected compact groups, like the orthogonal groups, that we would like to allow.

That is perhaps enough in the way of excuses and apologies to cover almost any definition. The one we will use is suggested by [**Kna**].

Definition 2.5. Write $GL(n)$ for the group of real or complex $n \times n$ matrices. The *Cartan involution* of $GL(n)$ is the automorphism conjugate transpose inverse:

$$\theta(g) = {}^t\overline{g}^{-1}.$$

A *linear reductive group* is a closed subgroup G of some $GL(n)$, preserved by θ and having finitely many connected components. A *reductive group* is a Lie group $\widetilde{G}$ endowed with a homomorphism $\pi\colon \widetilde{G} \to G$ onto a linear reductive group, so that the kernel of π is finite.

The group of fixed points of θ on $GL(n)$ is the orthogonal group $O(n)$ (in the case of $\mathbb{R}$) or the unitary group $U(n)$ (in the case of $\mathbb{C}$). The easiest examples of linear reductive groups are the closed subgroups of $O(n)$ or $U(n)$. All compact Lie groups are therefore linear reductive groups. Some additional examples are provided by Exercise 4.

The great advantage of Definition 2.5 is that it makes the most important structure theorem fairly easy to prove. Here it is.

Theorem 2.6 (Cartan decomposition)**.** *Suppose G is a linear reductive group. Write*

$$K = G^\theta, \qquad \mathfrak{s} = -1 \text{ eigenspace of } \theta \text{ on } \mathfrak{g}.$$

Then the map

$$K \times \mathfrak{s} \to G, \qquad (k, X) \mapsto k \exp(X)$$

is a diffeomorphism from $K \times \mathfrak{s}$ onto G. In particular K is maximal among the compact subgroups of G.

Suppose $\widetilde{G}$ is a reductive group, with $\pi\colon \widetilde{G} \to G$ as in Definition 2.5. Write $\widetilde{K} = \pi^{-1}(K)$, a compact subgroup of $\widetilde{G}$, and use $d\pi$ to identify the Lie algebras of $\widetilde{G}$ and G. Then the map

$$\widetilde{K} \times \mathfrak{s} \to \widetilde{G}, \qquad (\tilde{k}, X) \mapsto \tilde{k}\exp(X)$$

is a diffeomorphism from $\widetilde{K} \times \mathfrak{s}$ onto $\widetilde{G}$. In particular $\widetilde{K}$ is maximal among the compact subgroups of $\widetilde{G}$.

Suppose $\widetilde{G}$ is a reductive group. Define a map θ from $\widetilde{G}$ to itself by

$$\theta(\tilde{k}\exp(X)) = \tilde{k}\exp(-X) \qquad (\tilde{k} \in \widetilde{K}, X \in \mathfrak{s}).$$

Then θ is an involution of order 2, the Cartan involution of $\widetilde{G}$. The group of fixed points is $\widetilde{K}$.

For the case of $GL(n)$ this theorem is just the polar decomposition for matrices. The version for linear reductive groups follows fairly easily. The version for general reductive groups is a consequence.

The Lie algebra version of Theorem 2.6 is entirely trivial, but is worth stating explicitly. Notice that the Lie algebra $\mathfrak{k}$ of K is equal to the $+1$ eigenspace of the (involutive linear) automorphism θ of $\mathfrak{g}$. Since $\theta^2 = 1$, the Lie algebra must be the direct sum of the $+1$ and -1 eigenspaces of θ:

$$\mathfrak{g} = \mathfrak{k} + \mathfrak{s}.$$

It is often helpful to be able to identify the Lie algebra of a reductive group with its dual space. This we accomplish by means of the next lemma.

Proposition 2.7. *Suppose G is a reductive group. Identify $\mathfrak{g}$ with a Lie algebra of $n \times n$ matrices in accordance with Definition 2.5. Define a real-valued symmetric bilinear form on $\mathfrak{g}$ by*

$$\langle X, Y \rangle = \operatorname{Re}\operatorname{tr} XY.$$

1. *The form is invariant under $\operatorname{Ad}(G)$ and the Cartan involution θ:*
$$\langle \operatorname{Ad}(g)X, \operatorname{Ad}(g)Y \rangle = \langle X, Y \rangle = \langle \theta X, \theta Y \rangle.$$
2. *The Cartan decomposition $\mathfrak{g} = \mathfrak{k} + \mathfrak{s}$ is orthogonal for the form.*
3. *The form is positive definite on $\mathfrak{s}$ and negative definite on $\mathfrak{k}$, and (as a consequence) non-degenerate on $\mathfrak{g}$.*
4. *There is a G-equivariant linear isomorphism*
$$\mathfrak{g}^* \simeq \mathfrak{g}, \qquad \lambda \mapsto X_\lambda$$
characterized by
$$\lambda(Y) = \langle X_\lambda, Y \rangle \qquad (Y \in \mathfrak{g}).$$

The assertions about invariance are elementary, and the orthogonality follows from θ-invariance. The assertions about positivity can be proved easily for $GL(n)$ by explicit calculation with standard bases for matrices; the general cases follow. The last assertion is an easy formal consequence.

It follows from Proposition 2.7 that for G reductive, G orbits on $\mathfrak{g}^*$ may be identified with G orbits on $\mathfrak{g}$. In addition, these latter orbits are related to (more precisely, they are subsets of) conjugacy classes of matrices under $GL(n)$. The following definitions for matrices therefore make sense for $\mathfrak{g}$.

Definition 2.8. Suppose G is a reductive Lie group, with Lie algebra $\mathfrak{g}$ consisting of $n \times n$ matrices (Definition 2.5). An element $X \in \mathfrak{g}$ is called *nilpotent* if it is nilpotent as a matrix; that is, if $X^N = 0$ for N large enough. An equivalent requirement is that every eigenvalue of X be equal to zero.

An element $X \in \mathfrak{g}$ is called *semisimple* if the corresponding complex matrix is diagonalizable. (By this we mean that X is regarded as an $n \times n$ complex matrix, even if $\mathfrak{g}$ consists of real matrices. For example, the real matrix

$$X = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

is semisimple, since over $\mathbb{C}$ it is conjugate to $\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$.)

An element $X \in \mathfrak{g}$ is called *hyperbolic* if it is semisimple and all eigenvalues are real. Typical examples are the self-adjoint matrices in $\mathfrak{s}$.

An element $X \in \mathfrak{g}$ is called *elliptic* if it is semisimple and all eigenvalues are purely imaginary. Typical examples are the skew-adjoint matrices in $\mathfrak{k}$, like the 2×2 example above.

It is more difficult to produce examples of nilpotent elements related to the Cartan involution and the reductive group structure. Here is one way. Suppose G is a linear reductive group, and we have three elements H, E, and F of $\mathfrak{g}$, satisfying

$$\theta E = -F, \quad \theta H = -H, \quad [H, E] = 2E, \quad [E, F] = H. \tag{2.9(a)}$$

Then necessarily E is nilpotent; for the matrix E^N belongs to the $2N$ eigenspace of $\mathrm{ad}(H)$ on $\mathfrak{gl}(n)$, and this must be zero for large N. On the other hand, there is a group homomorphism

$$\phi\colon SL(2, \mathbb{R}) \to G, \tag{2.9(b)}$$

characterized by the requirements

$$d\phi \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = H, \quad d\phi \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = E, \quad d\phi \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} = F. \tag{2.9(c)}$$

This homomorphism respects the Cartan involutions:

$$\phi({}^t g^{-1}) = \theta(\phi(g)). \tag{2.9(d)}$$

Exercise 6 asks you to think about why the homomorphism ϕ exists.

Evidently a homomorphism ϕ satisfying (2.9)(d) defines by (2.9)(c) elements H, E, and F satisfying (2.9)(a). We call such elements a *standard* $\mathfrak{sl}(2)$ *triple.* The element E alone determines $F = -\theta E$ and $H = [E, F]$. Not every nilpotent element E belongs to a standard triple, however.

Proposition 2.10 (Jordan decomposition)**.** *Suppose G is a reductive group with Lie algebra $\mathfrak{g}$ (Definition 2.5) and Cartan decomposition $G = K \cdot \exp(\mathfrak{s})$ (Theorem 2.6).*

1. *Any element $X \in \mathfrak{g}$ has a unique decomposition*

$$X = X_h + X_e + X_n$$

characterized by the requirements that X_h is hyperbolic (Definition 2.8); X_e is elliptic; X_n is nilpotent; and X_h, X_e, and X_n commute with each other.

2. *After replacing X by a conjugate under $\mathrm{Ad}(G)$, we may assume that $X_h \in \mathfrak{s}$, $X_e \in \mathfrak{k}$, and that $X_n = E$ belongs to a standard $\mathfrak{sl}(2)$ triple (see (2.9)(a)).*
3. *The $\mathrm{Ad}(G)$ orbits of hyperbolic elements in $\mathfrak{g}$ are in one-to-one correspondence with the $\mathrm{Ad}(K)$ orbits on $\mathfrak{s}$.*
4. *The $\mathrm{Ad}(G)$ orbits of elliptic elements in $\mathfrak{g}$ are in one-to-one correspondence with the $\mathrm{Ad}(K)$ orbits on $\mathfrak{k}$.*
5. *The $\mathrm{Ad}(G)$ orbits of nilpotent elements in $\mathfrak{g}$ are in one-to-one correspondence with the $\mathrm{Ad}(K)$ orbits of standard $\mathfrak{sl}(2)$ triples in $\mathfrak{g}$.*

Proposition 2.10 provides a systematic way to list the orbits of a reductive group G acting on its Lie algebra. First, list the possible hyperbolic parts X_h. According to the proposition, this is equivalent to listing the orbits of K on $\mathfrak{s}$. (By standard structure theory, this in turn is equivalent to listing the orbits of the "small Weyl group" of G on a Cartan subspace $\mathfrak{a}$; but we will not use this fact.)

Proposition 2.11. *Suppose $G = K \cdot \exp(\mathfrak{s})$ is the Cartan decomposition of a reductive group, and $X_h \in \mathfrak{s}$. Write G^{X_h} for the stabilizer of X_h in the adjoint action. Then G^{X_h} is a reductive group with Cartan decomposition $K^{X_h} \cdot \exp(\mathfrak{s}^{X_h})$. The orbits of G on $\mathfrak{g}$ with hyperbolic part X_h are in one-to-one correspondence with the orbits of G^{X_h} on $\mathfrak{g}^{X_h}$ having hyperbolic part 0. The correspondence sends the orbit of $X_e + X_n$ for G^{X_h} to the orbit of $X_h + X_e + X_n$.*

The only part of this proposition that does not follow at once from Proposition 2.10 is the fact that G^{X_h} is a reductive group. This may be reduced at once to the case that G is linear reductive. Then the fact that G^{X_h} is closed and θ-stable is easy; what requires proof is the finiteness of the group of connected components. There are some hints in Exercise 7.

So now we are reduced to the problem of listing orbits with hyperbolic part 0, in a smaller reductive group. To keep the notation simple, we will just discuss this problem for G itself. We proceed by listing all possible elliptic parts X_e for such an orbit. According to Proposition 2.10, this is the same as listing the orbits of K on its own Lie algebra $\mathfrak{k}$. (Again standard structure theory reduces this to the orbits of the Weyl group of K on the Lie algebra of a maximal torus, but again we do not need this.)

Proposition 2.12. *Suppose $G = K \cdot \exp(\mathfrak{s})$ is the Cartan decomposition of a reductive group, and $X_e \in \mathfrak{k}$. Write G^{X_e} for the stabilizer of X_e in the adjoint action. Then G^{X_e} is a reductive group with Cartan decomposition $K^{X_e} \cdot \exp(\mathfrak{s}^{X_e})$. The orbits of G on $\mathfrak{g}$ with hyperbolic part 0 and elliptic part X_e are in one-to-one correspondence with the nilpotent orbits of G^{X_e} on $\mathfrak{g}^{X_e}$. The correspondence sends the orbit of X_n for G^{X_e} to the orbit of $X_e + X_n$.*

This is an easy consequence of Proposition 2.10.

Finally we are reduced to listing nilpotent orbits in a still smaller reductive group. (It is worth remarking that the class of smaller reductive groups arising in these reductions from a fixed G is finite.) In order to complete the picture in Propositions 2.11 and 2.12, we restate the description of nilpotent orbits in Proposition 2.10.

Proposition 2.13. *Suppose G is a linear reductive group. The nilpotent orbits of G on $\mathfrak{g}$ are in one-to-one correspondence with the K-conjugacy classes of homomorphisms*

$$\phi\colon SL(2,\mathbb{R}) \to G$$

respecting the Cartan involutions: $\phi({}^t g^{-1}) = \theta(\phi(g))$. The correspondence sends ϕ to the orbit of $d\phi\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$.

For any reductive G there is the trivial homomorphism ϕ, corresponding to the nilpotent orbit 0. There will be other classes if and only if the derived group of G_0 is noncompact. In any case the total number of classes (the number of nilpotent orbits in the Lie algebra) is finite. This is not particularly obvious or easy to prove; a nice introduction to the theory is in [**CM**].

Recall that we are really concerned with representation theory. Our interest in orbits arises from the formal relationship suggested by mathematical physics between Definitions 1.8 and 1.9: that there should be a method of "quantization" to pass from coadjoint orbits for G to irreducible unitary representations of G. According to Proposition 2.7, we can identify coadjoint orbits with adjoint orbits (for reductive groups). The last few propositions have described the adjoint orbits in three steps. The idea of quantization correspondingly suggests that unitary representations of a reductive group G should be constructed in three steps; even more, that an irreducible unitary representation ought to have a kind of "Jordan decomposition" along the lines of Proposition 2.10. We will not consider this problem very seriously, but here is a brief outline of what is known.

The first step in the Jordan decomposition is finding the hyperbolic part of an orbit. For this step there is a perfect and complete analogue in unitary representation theory, given in [**Kna**], Theorem 16.10. The second step is finding the elliptic part of an orbit. Here the representation-theoretic analogue is known (see [**VoCI**]) but its behavior is more complicated. A partly conjectural connection with the behavior of orbits is described in [**SV**]; the geometric nature of the results is emphasized in [**VLC**].

The final step in the Jordan decomposition is a nilpotent coadjoint orbit. The representation-theoretic analogue is the theory of unipotent representations, which is still in its infancy. Two more or less expository discussions of the theory may be found in [**VO**] and [**VAV**]. As indicated in the introduction, we will return to the problem of quantizing nilpotent coadjoint orbits in the latter part of these lectures.

The title of this section suggests that the three steps for constructing a unitary representation from a coadjoint orbit should be different from each other, even though they share a common motivation. In a sense accepting such differences is an admission of defeat, and contrary to the spirit of the orbit method. Nevertheless it is of tremendous value technically. To conclude this section, we will outline the construction of representations parallel to the hyperbolic part of a coadjoint orbit: parabolic induction. This construction appeared in a number of special cases before Gelfand and his collaborators began to emphasize its general importance in the early 1950s. The history of the subject from that point on is colorful and impressive, but my scholarship is not deep enough, nor my hide thick enough, to allow a discussion of it here.

Let us therefore fix a hyperbolic element

$$\lambda_h \in \mathfrak{g}^* \tag{2.14)(a}$$

in the dual of the Lie algebra for our reductive Lie group G. Recall from Proposition 2.7 that this means that the element $X_h \in \mathfrak{g}$ defined by

$$\lambda_h(Y) = \langle X_h, Y \rangle \qquad (Y \in \mathfrak{g}) \tag{2.14(b)}$$

is hyperbolic. After replacing λ_h by a conjugate under G, we may therefore assume that

$$X_h \in \mathfrak{s}. \tag{2.14(c)}$$

The isotropy group for the coadjoint action at λ_h is

$$G^{\lambda_h} = \{g \in G \mid \mathrm{Ad}(g)(X_h) = X_h\} = L, \tag{2.14(d)}$$

a reductive group with Cartan involution $\theta|_L$ (Proposition 2.11). Now a hyperbolic element of the Lie algebra acts in any algebraic representation by a diagonalizable operator with real eigenvalues. (Exercise 8 outlines a proof of this fact for the adjoint representation.) Consequently $\mathrm{ad}(X_h)$ is diagonalizable with real eigenvalues. Explicitly,

$$\mathfrak{g} = \sum_{r \in \mathbb{R}} \mathfrak{g}_r, \tag{2.14(e)}$$

with

$$\mathfrak{g}_r = \{Y \in \mathfrak{g} \mid [X_h, Y] = rY\}. \tag{2.14(f)}$$

Two immediate consequences are

$$[\mathfrak{g}_r, \mathfrak{g}_s] \subset \mathfrak{g}_{r+s}, \qquad \mathrm{Ad}(L)(\mathfrak{g}_r) = \mathfrak{g}_r, \qquad \mathfrak{g}_0 = \mathrm{Lie}(L). \tag{2.14(g)}$$

Proposition 2.15. *With notation as above, define*

$$\mathfrak{u} = \sum_{r > 0} \mathfrak{g}_r. \tag{2.15(a)}$$

Then $\mathfrak{u}$ is an ad*-nilpotent Lie subalgebra of $\mathfrak{g}$, normalized by the adjoint action of L. The exponential map is a diffeomorphism of $\mathfrak{u}$ onto a closed subgroup U of G, also normalized by L and meeting it trivially. The semidirect product group $Q = LU \subset G$ is a closed subgroup of G. The homogeneous space G/Q is compact; in fact it is homeomorphic to $K/L \cap K$.*

There is a unitary character $\chi(\lambda_h)$ of L defined by

$$\chi(\lambda_h)(k \cdot \exp(Z)) = \exp(i\lambda_h(Z)) \qquad (k \in L \cap K, Z \in \mathfrak{l} \cap \mathfrak{s}). \tag{2.15(b)}$$

Equivalently, $\chi(\lambda_h)$ may be defined as the unique character of L that is trivial on $L \cap K$, and has differential $i\lambda_h$.

There is a character ρ_Q of Q taking positive real values, defined by

$$\rho_Q(q) = |\det(\mathrm{Ad}(q)|_{\mathfrak{u}})|^{1/2}. \tag{2.15(c)}$$

Many of the assertions in this proposition are very easy to prove, and none is particularly deep. We will not discuss the proof, however.

The character ρ_Q is included because it is used in the definition of Mackey's unitary induction from Q to G. Geometrically, it is the character of Q corresponding to the (Q-equivariant) bundle of half-densities on G/Q. To see that, it is better to write

$$\rho_Q(q) = |\det\left(\mathrm{Ad}^*(q)|_{(\mathfrak{g}/\mathfrak{q})^*}\right)|^{1/2}.$$

But the equivalence of this definition with (2.15)(c) is easy to prove.

Before describing how this structure is used to construct unitary representations, we will say a little about its geometric content. We should look at the geometry of arbitrary coadjoint orbits with hyperbolic part λ_h, but to simplify the discussion we consider only the coadjoint orbit of λ_h itself. This is

$$X = \mathrm{Ad}^*(G)\cdot\lambda_h \simeq G/G^{\lambda_h} = G/L. \tag{2.16(a)}$$

Because L is contained in Q, there is a natural fibration

$$X \simeq G/L \to G/Q; \tag{2.16(b)}$$

the fiber over the identity coset is $Q/L \simeq U$. Using the definition of the symplectic structure on X in Lemma 2.1, together with the structural information in (2.15), one sees easily that Q/L is a Lagrangian submanifold of X; that is, that Q/L has half the dimension of $X = G/L$, and that the symplectic form ω_λ vanishes on $\mathfrak{q}/\mathfrak{l}$. We summarize this by saying that

$$X \simeq G/L \to G/Q \text{ is a } G\text{-equivariant Lagrangian fibration.} \tag{2.16(c)}$$

Definition 2.17. Suppose G is a reductive Lie group, λ_h is a hyperbolic element of $\mathfrak{g}^*$ as in (2.14), L is the stabilizer of λ_h in G, and $Q = LU$ is the semidirect product subgroup constructed in Proposition 2.15. The functor of *parabolic induction* carries unitary representations $(\pi_L, \mathcal{H}_L)$ of L to unitary representations $(\pi_G, \mathcal{H}_G)$ of G, as follows. First extend π_L to a representation π_Q of Q on the same Hilbert space, by making U act trivially. Next, define $\mathcal{H}_G^c$ to be the vector space of continuous functions from G to $\mathcal{H}_L$ satisfying

$$\phi(gq) = \rho_Q(q^{-1})\pi_Q(q^{-1})\phi(g) \qquad (g\in G, q\in Q). \tag{2.17(a)}$$

The group G acts on $\mathcal{H}_G^c$ by left translation:

$$[\pi_G^c(g)\phi](x) = \phi(g^{-1}x). \tag{2.17(b)}$$

We make $\mathcal{H}_G^c$ into a pre-Hilbert space as follows. Suppose ϕ and ψ belong to $\mathcal{H}_G^c$. Define a complex-valued function $f_{\langle\phi,\psi\rangle}$ on G by

$$f_{\langle\phi,\psi\rangle}(g) = \langle\phi(g), \psi(g)\rangle; \tag{2.17(c)}$$

the inner product on the right is in $\mathcal{H}_L$. From the transformation property (2.17)(a) and the unitarity of π_Q we deduce

$$f_{\langle\phi,\psi\rangle}(gq) = \rho_Q^2(q^{-1})f_{\langle\phi,\psi\rangle}(g). \tag{2.17(d)}$$

Up to a choice of measure on $\mathfrak{g}/\mathfrak{q}$, this means that $f_{\langle\phi,\psi\rangle}$ may be regarded as a continuous density on G/Q. Since G/Q is compact, such a density has finite total volume; so we can define

$$\langle\phi,\psi\rangle = \int_{G/Q} f_{\langle\phi,\psi\rangle}. \tag{2.17(e)}$$

That this is a pre-Hilbert structure on $\mathcal{H}_G^c$ is clear. Define $\mathcal{H}_G$ to be the completion of $\mathcal{H}_G^c$. Since the operators $\pi_G^c(g)$ preserve the pre-Hilbert structure, they extend to unitary operators $\pi_G(g)$ on $\mathcal{H}_G$. It is straightforward to show that

$$(\pi_G, \mathcal{H}_G) = \mathrm{Ind}_Q^G(\pi_L \otimes 1) \qquad (2.17)(f)$$

is a (continuous) unitary representation of G.

The simplest example of Definition 2.17 has π_L equal to the trivial representation of L. In that case π_G is the representation of G on the space of square-integrable half-densities on G/Q. Accordingly to the philosophy of geometric quantization, the corresponding Hamiltonian G-space (in the correspondence we seek between Definitions 1.8 and 1.9) is just the cotangent bundle $T^*(G/Q)$. This space is not homogeneous for G; the moment map to $\mathfrak{g}^*$ has its image in the cone of nilpotent elements. (This means that Definition 2.17 has something to teach us about quantizing nilpotent coadjoint orbits. We will not pursue this knowledge, however.)

The next simplest example of Definition 2.17 has π_L equal to the one-dimensional character $\chi(\lambda_h)$ of Proposition 2.15. In that case π_G is a space of sections of the bundle of half-densities on G/Q, twisted by a Hermitian line bundle defined by $\chi(\lambda_h)$. It should be thought of as a quantization of the hyperbolic coadjoint orbit $X = \mathrm{Ad}^*(G) \cdot \lambda_h$.

Parabolic induction turns out to be a nice analogue for unitary representations of the geometric correspondence described in Proposition 2.11. There is not really a theorem to state here, but we can make a definition.

Definition 2.18. Suppose G is a reductive Lie group, $\lambda_h \in \mathfrak{g}^*$ is a hyperbolic element, L is the stabilizer of λ_h in G, and $Q = LU$ is the group constructed in Proposition 2.15. Recall from Proposition 2.11 that there is a bijection between coadjoint orbits for G with hyperbolic part conjugate to λ_h, and coadjoint orbits for L with hyperbolic part 0.

Let $X_L \subset \mathfrak{l}^*$ be any coadjoint orbit for L with hyperbolic part 0, and suppose π_L is any unitary representation of L attached to X_L. (Here "attached" refers to some unspecified quantization scheme for L; it does not need to have any particular mathematical meaning.) Write $X_G = \mathrm{Ad}^*(G) \cdot (\lambda_h + X_L)$ for the corresponding coadjoint orbit for G. Then we declare that the unitary representation of G

$$\pi_G = \mathrm{Ind}_Q^G(\pi_L \otimes \chi(\lambda_h))$$

(Definition 2.17) is attached to X_G.

If we want to make the orbit method into a scheme for classifying unitary representations, then there are theorems to be proven here: that (for appropriate classes of π_L) the representations π_G are irreducible, and that they exhaust the irreducible unitary representations of G. That is the nature of the results in [**Kna**], Theorem 16.10 (compare also the formulation in [**VLC**], section 3). Since we are concerned mostly with defining a method of quantization, we will be content with Definitions 2.17 and 2.18.

LECTURE 3
Complex polarizations

Our goal in this section is to introduce briefly the replacement for parabolic induction related to elliptic coadjoint orbits. Because these ideas will be at the center of the lectures of Roger Zierau, we will give few details. We want only to formulate the main ideas in such a way as to suggest a transition from the hyperbolic case to the nilpotent case.

We begin by considering again the geometric content of parabolic induction; specifically, how the parabolically induced representation can be described in terms of the coadjoint orbit to which it is formally attached. To simplify, we will work only with the hyperbolic orbit $X = G \cdot \lambda_h \simeq G/L$, and not with a more general orbit having hyperbolic part λ_h. We first defined a unitary character $\chi(\lambda_h)$ of L, and thus a hermitian line bundle on G/L. Next, we twisted this line bundle by a half-density bundle attached to a character ρ_Q of L. Finally, we considered sections of the twisted line bundle on G/Q; that is, sections on G/L that are constant along the fibers of the projection $G/L \to G/Q$. Because these fibers are connected, the constancy can be expressed in terms of the vanishing of certain derivatives. Here is a slightly more precise statement.

Proposition 3.1. *In the setting of Definition 2.17, the space of smooth vectors in the representation* $\operatorname{Ind}_Q^G(\chi(\lambda_h))$ *may be identified with smooth functions* ϕ *on* G *satisfying*

$$\phi(gl) = \rho_Q(l^{-1})\chi(\lambda_h)(l^{-1})\phi(g) \qquad (g \in G, l \in L),$$

and

$$Z \cdot \phi = 0 \qquad (Z \in \mathfrak{u}).$$

Here in the second condition we identify $\mathfrak{g}$ *with left-invariant vector fields on* G*; that is, with right derivatives.*

The advantage of this formulation over Definition 2.17 is the almost complete disappearance of the group Q (which has no analogue in the case of elliptic coadjoint orbits).

We begin now a discussion of quantizing elliptic coadjoint orbits that is parallel to that for hyperbolic orbits in (2.14). We fix therefore an elliptic element

$$\lambda_e \in \mathfrak{g}^*, \tag{3.2)(a}$$

and write $X_e \in \mathfrak{g}$ for its image under the isomorphism of Proposition 2.7. After conjugating λ_e by G, we may assume that $\theta\lambda_e = \lambda_e$, or equivalently that $X_e \in \mathfrak{k}$.

The isotropy group for the coadjoint action at λ_e is

$$G^{\lambda_e} = \{g \in G \mid \mathrm{Ad}(g)(X_e) = X_e\} = L, \tag{3.2)(b}$$

a closed reductive subgroup of G with Cartan involution $\theta|_L$. Elements of the Lie algebra of the compact group K must act in any finite-dimensional representation of K by diagonalizable operators with purely imaginary eigenvalues. Consequently $i\,\mathrm{ad}(X_e)_{\mathbb{C}}$ is diagonalizable with real eigenvalues. Explicitly,

$$\mathfrak{g}_{\mathbb{C}} = \sum_{r\in\mathbb{R}} \mathfrak{g}_{\mathbb{C},r}, \tag{3.2)(c}$$

with

$$\mathfrak{g}_{\mathbb{C},r} = \{Y \in \mathfrak{g}_{\mathbb{C}} \mid [iX_e, Y] = rY\}. \tag{3.2)(d}$$

Just as in (2.14), we have

$$[\mathfrak{g}_{\mathbb{C},r}, \mathfrak{g}_{\mathbb{C},s}] \subset \mathfrak{g}_{\mathbb{C},r+s}, \qquad \mathrm{Ad}(L)(\mathfrak{g}_{\mathbb{C},r}) = \mathfrak{g}_{\mathbb{C},r}, \qquad \mathfrak{g}_{\mathbb{C},0} = \mathrm{Lie}(L)_{\mathbb{C}}. \tag{3.2)(e}$$

Just as in Proposition 2.15, we can define

$$\mathfrak{u} = \sum_{r>0} \mathfrak{g}_{\mathbb{C},r}. \tag{3.2)(f}$$

Then $\mathfrak{u}$ is an ad-nilpotent Lie subalgebra of $\mathfrak{g}_{\mathbb{C}}$, normalized by the adjoint action of L. The semidirect product Lie algebra

$$\mathfrak{q} = \mathfrak{l}_{\mathbb{C}} + \mathfrak{u} \subset \mathfrak{g}_{\mathbb{C}} \tag{3.2)(g}$$

is a parabolic subalgebra of $\mathfrak{g}_{\mathbb{C}}$. What is different from the hyperbolic case is that these are not complexifications of subalgebras of $\mathfrak{g}$. To see that more precisely, we consider complex conjugation acting on $\mathfrak{g}_{\mathbb{C}}$:

$$\overline{U + iV} = U - iV \qquad (U, V \in \mathfrak{g}).$$

Then $\overline{\mathfrak{g}_{\mathbb{C},r}} = \mathfrak{g}_{\mathbb{C},-r}$. Therefore

$$\overline{\mathfrak{u}} = \sum_{r>0} \mathfrak{g}_{\mathbb{C},-r}, \quad \overline{\mathfrak{q}} = \mathfrak{l}_{\mathbb{C}} + \overline{\mathfrak{u}}. \tag{3.2)(h}$$

It follows that

$$\mathfrak{q} + \overline{\mathfrak{q}} = \mathfrak{g}_{\mathbb{C}}, \quad \mathfrak{q} \cap \overline{\mathfrak{q}} = \mathfrak{l}_{\mathbb{C}}, \quad \mathfrak{g}_{\mathbb{C}} = \overline{\mathfrak{u}} \oplus \mathfrak{l}_{\mathbb{C}} \oplus \mathfrak{u}. \tag{3.2)(i}$$

Proposition 3.3. *Suppose $X = G \cdot \lambda_e \simeq G/L$ is an elliptic coadjoint orbit for the reductive group G; use the notation above. Then there is a distinguished G-invariant complex structure on X, characterized by the requirement that $\mathfrak{q}/\mathfrak{l}_{\mathbb{C}} \subset \mathfrak{g}_{\mathbb{C}}/\mathfrak{l}_{\mathbb{C}}$ is the antiholomorphic tangent space to X at the point λ_e (corresponding to the identity coset eL).*

There is a unitary character $2\rho_{\mathfrak{q}}$ of L defined by

$$2\rho_{\mathfrak{q}}(l) = \det(\mathrm{Ad}(l)|_{\mathfrak{u}}). \tag{3.3)(a}$$

An equivalent definition is

$$2\rho_{\mathfrak{q}}(l) = \det(\mathrm{Ad}^*(l)|_{(\mathfrak{g}/\mathfrak{q})^*}). \tag{3.3(b)}$$

The differential of $2\rho_{\mathfrak{q}}$ is a purely imaginary-valued one-dimensional character of the Lie algebra $\mathfrak{l}$, which we still denote by $2\rho_{\mathfrak{q}}$. Explicitly,

$$2\rho_{\mathfrak{q}}(Z) = \operatorname{tr}\operatorname{ad}(Z)|_{\mathfrak{u}} \qquad (Z \in \mathfrak{l}). \tag{3.3(c)}$$

The linear functional $i\lambda_e$ restricted to $\mathfrak{l}$ is a purely imaginary-valued one-dimensional character of the Lie algebra.

We now have in hand most of the ingredients we need to formulate a definition along the lines of Proposition 3.1 for elliptic orbits. There are two problems. First, instead of the character ρ_Q of L we have available only $2\rho_{\mathfrak{q}}$, which is like the square of ρ_Q. We defined ρ_Q by taking a square root (of an absolute value, losing some plus or minus one factors in the process). The character $2\rho_{\mathfrak{q}}$ is unitary, so its absolute value is identically one (and uninteresting). If we do not take an absolute value, the square root may or may not exist (as a unitary character of L). (Exercise 9 looks at an example.)

The second problem is that the linear functional $i\lambda_e$ may not exponentiate to a unitary character of L. This problem is of fundamental importance: it's not a bug, it's a feature. The "elliptic unitary representations" parametrized by elliptic coadjoint orbits come not in continuous families (as the orbits do, by Proposition 2.10(4)) but in discrete ones. For representation theory we are interested only in the elliptic orbits to which we can attach something like a unitary character of L. The precise requirement addresses both problems (the square root of $2\rho_{\mathfrak{q}}$ and the exponential of $i\lambda_e$) at the same time.

Definition 3.4. Suppose G is a reductive Lie group, and $\lambda_e \in \mathfrak{g}^*$ is an elliptic element. Use the notation of (3.2) and (3.3). An *admissible datum at λ_e* is an irreducible unitary representation (τ, V_τ) of L so that

$$d\tau = (i\lambda_e + \rho_{\mathfrak{q}}) \cdot I_\tau;$$

here I_τ is the identity operator on the vector space V_τ.

The representation V_τ gives rise to a G-equivariant vector bundle $\mathcal{V}_\tau$ over G/L; the smooth sections of $\mathcal{V}_\tau$ may be identified with smooth functions ϕ from G to V_τ satisfying

$$\phi(gl) = \tau(l^{-1})\phi(g) \qquad (g \in G, l \in L). \tag{3.4(a)}$$

This identification makes sense locally on G/L: if U is an open set in G/L and $\widetilde{U}$ its preimage in G, then smooth sections of $\mathcal{V}_\tau$ over U correspond to functions on $\widetilde{U}$ satisfying (3.4)(a) for all $g \in \widetilde{U}$.

Recall now the holomorphic structure on G/L introduced in Proposition 3.3. The vector bundle $\mathcal{V}_\tau$ is naturally a G-equivariant holomorphic bundle. The holomorphic sections of $\mathcal{V}_\tau$ over an open set U may be identified with smooth functions ϕ from $\widetilde{U}$ to V_τ satisfying

$$\phi(gl) = \tau(l^{-1})\phi(g) \qquad (g \in \widetilde{U}, l \in L), \tag{3.4(b)}$$

and

$$Z \cdot \phi = 0 \qquad (Z \in \mathfrak{u}). \tag{3.4)(c}$$

Here in the second condition we identify $\mathfrak{g}_{\mathbb{C}}$ with left-invariant complex vector fields on G; these are just the Cauchy-Riemann equations for the corresponding section. The group G acts by left translation on holomorphic sections of $\mathcal{V}_\tau$.

This definition of admissible orbit datum does not look much like Duflo's (see [**D**] or [**VAV**], Definition 7.2). Nevertheless it is equivalent (for elliptic orbits in reductive groups). A reformulation that is a little closer to Duflo's definition appears at the end of this section.

Given λ_e, there are just finitely many possibilities for τ (possibly none). All of them are finite-dimensional. If G is connected, then so is L, so there is at most one possible τ in that case.

Definition 3.4 provides us with a natural analogue of the representation attached to the hyperbolic orbit, as described in Proposition 3.1. It is the space $\Gamma(G/L, \mathcal{V}_\tau)$ of all holomorphic sections of $\mathcal{V}_\tau$. This is a reasonable topological vector space, endowed with a natural continuous representation of G (by left translation). There are problems, however. One is that G/L is usually not a Stein manifold, so there is no reason to expect random holomorphic vector bundles to have any global sections at all. Indeed $\mathcal{V}_\tau$ rarely has holomorphic sections (see Exercise 10).

In the absence of holomorphic sections, it is natural to look at higher Dolbeault cohomology of G/L with coefficients in $\mathcal{V}_\tau$. In many cases this leads to nice topological representations of G, but it is always difficult (and often impossible) to find G-invariant pre-Hilbert space structures on these representations. The problem appears even when G/L is Stein. The natural pre-Hilbert space structure is given by integrating the Hermitian inner product on sections over G/L. This integral does not converge for all holomorphic sections, and in some interesting cases it does not converge except for the zero section. For the representations on higher Dolbeault cohomology the problems are worse, and they are far from completely understood; Zierau's lectures will explain what is known.

For the purpose of abstract quantization—of formally attaching unitary representations to coadjoint orbits—many of these problems can be circumvented. This is the subject of [**KV**], which I will not attempt to summarize here. There is in particular a definition like Definition 2.18, saying how to attach unitary representations to "elliptic plus nilpotent" coadjoint orbits once we know how to attach them to nilpotent orbits. An exposition appears in [**VUR**]. We will be content with the following special case.

Proposition 3.5 (see [**Wng**]). *Suppose G is a reductive Lie group, and $\lambda_e \in \mathfrak{g}^*$ is an elliptic element. Use the notation of Definition 3.4. Suppose that (τ, V_τ) is an admissible datum at λ_e. Then there is attached to the G-orbit of (λ_e, τ) a unitary representation $(\pi(\lambda_e, \tau), \mathcal{H})$ of G, characterized by the following properties. Write $\mathcal{V}_\tau$ for the holomorphic vector bundle on $X = G/L$ associated to τ, and $s = \dim_{\mathbb{C}}(K/L \cap K)$.*

1. *The Dolbeault cohomology groups $H^{0,p}(X, \mathcal{V}_\tau)$ vanish for $p \neq s$.*
2. *The $\overline{\partial}$ operator computing Dolbeault cohomology has closed range. Consequently the cohomology space $H^{0,s}(X, \mathcal{V}_\tau)$ inherits a nice complete Hausdorff topological vector space structure, on which G acts continuously by left translation.*

3. *There is a G-equivariant embedding with dense image of the Hilbert space representation $\mathcal{H}$ in $H^{0,s}(X, \mathcal{V}_\tau)$.*

The representation $\pi(\lambda_e, \tau)$ is always a finite sum of irreducible representations; but it may vanish, and it may be reducible. The second possibility in particular sounds at first like a bug in the orbit method. For an attempt to make it a feature, see [**VDK**], Theorem 10.12(e).

To conclude this section, we recast the definition of admissible datum (Definition 3.4). We first construct a square root of $2\rho_\mathfrak{q}$. For that, define the *metaplectic double cover of L* by

$$\widetilde{L} = \{(l, z) \in L \times \mathbb{C}^\times \mid 2\rho_\mathfrak{q}(l) = z^2\}. \tag{3.6(a)}$$

Projection on the first factor is a surjective homomorphism

$$p \colon \widetilde{L} \to L, \qquad p(l, z) = l. \tag{3.6(b)}$$

Evidently the kernel of p is the two-element central subgroup $\{1, \epsilon\}$, with $\epsilon = (1, -1) \in \widetilde{L}$. Projection on the *second* factor defines a unitary character $\rho_\mathfrak{q}$ of $\widetilde{L}$:

$$\rho_\mathfrak{q}(l, z) = z. \tag{3.6(c)}$$

The notation is chosen because the square of this character is just $2\rho_\mathfrak{q}$: explicitly,

$$[\rho_\mathfrak{q}(x)]^2 = 2\rho_\mathfrak{q}(p(x)) \qquad (x \in \widetilde{L}). \tag{3.6(d)}$$

We call a representation of $\widetilde{L}$ *genuine* if the central element ϵ acts by -1. A *metaplectic admissible datum at λ_e* is a genuine irreducible unitary representation $(\widetilde{\tau}, V_{\widetilde{\tau}})$ of $\widetilde{L}$, with the property that

$$d\widetilde{\tau} = i\lambda_e \cdot I_{\widetilde{\tau}}. \tag{3.6(e)}$$

Proposition 3.7. *Suppose $\lambda_e \in \mathfrak{g}^*$ is an elliptic element; use the notation of Definition 3.4 and (3.6). There is a one-to-one correspondence between admissible data (τ, V_τ) at λ_e, and metaplectic admissible data $(\widetilde{\tau}, V_{\widetilde{\tau}})$, characterized by*

$$\widetilde{\tau} \otimes \rho_\mathfrak{q} = \tau \circ p.$$

This is very easy. The new definition of admissible datum still differs from Duflo's in that Duflo's construction of the covering $\widetilde{L}$ is different; but the two constructions may be shown to be equivalent.

LECTURE 4
The Kostant-Sekiguchi correspondence

We have now seen more or less how to carry out the first two steps of the quantization program described in section 2. What remains is to quantize nilpotent coadjoint orbits. This we will not be able to do completely, but the effort will occupy the rest of the notes. Our first goal is to understand more precisely the structure and classification of these orbits.

We begin by sharpening and recasting Proposition 2.13. To begin, we introduce some notation for $SL(2,\mathbb{R})$ and its complexification $SL(2,\mathbb{C})$. Write θ_0 for the inverse transpose automorphism:

$$\theta_0(g) = {}^t g^{-1} \qquad (g \in SL(2,\mathbb{C})). \tag{4.1)(a}$$

We denote its differential by the same letter:

$$\theta_0(Z) = -{}^t Z \qquad (Z \in \mathfrak{sl}(2,\mathbb{C})). \tag{4.1)(b}$$

The complex conjugation σ_0 defining the real form $SL(2,\mathbb{R})$ is just complex conjugation of matrices.

The standard basis of $\mathfrak{sl}(2,\mathbb{R})$ is

$$H_0 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad E_0 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad F_0 = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}. \tag{4.1)(c}$$

These matrices satisfy

$$\begin{aligned} &[H_0, E_0] = 2E_0, \qquad [H_0, F_0] = -2F_0, \qquad [E_0, F_0] = H_0; && (4.1)(d) \\ &\theta_0(H_0) = -H_0, \qquad \theta_0(E_0) = -F_0, \qquad \theta_0(F_0) = -E_0. && (4.1)(e) \end{aligned}$$

All are fixed by σ_0. We will also need a different basis of $\mathfrak{sl}(2,\mathbb{C})$ diagonalizing the action of θ_0:

$$h_0 = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}, \quad x_0 = \frac{1}{2}\begin{pmatrix} 1 & -i \\ -i & -1 \end{pmatrix}, \quad y_0 = \frac{1}{2}\begin{pmatrix} 1 & i \\ i & -1 \end{pmatrix}. \tag{4.1)(f}$$

These matrices satisfy

$$\begin{aligned} &[h_0, x_0] = 2x_0, \qquad [h_0, y_0] = -2y_0, \qquad [x_0, y_0] = h_0; && (4.1)(g) \\ &\theta_0(h_0) = h_0, \qquad \theta_0(x_0) = -x_0, \qquad \theta_0(y_0) = -y_0; && (4.1)(h) \\ &\sigma_0(h_0) = -h_0, \qquad \sigma_0(x_0) = y_0, \qquad \sigma_0(x_0) = y_0. && (4.1)(i) \end{aligned}$$

Here is the improvement of Proposition 2.13 that we want.

Proposition 4.2. *Suppose G is a real reductive group with Cartan involution θ and corresponding maximal compact subgroup K. Write $\mathfrak{g}_{\mathbb{C}}$ and $K_{\mathbb{C}}$ for the complexifications of $\mathfrak{g}$ and K, and σ for the complex conjugation on $\mathfrak{g}_{\mathbb{C}}$. Then the following sets are in natural one-to-one correspondence.*

1. *Orbits of G on the cone $\mathcal{N}_{\mathbb{R}}$ of nilpotent elements of $\mathfrak{g}$.*
2. *Conjugacy classes under G of Lie algebra homomorphisms $\phi_{\mathbb{R}}$ from $\mathfrak{sl}(2,\mathbb{C})$ to $\mathfrak{g}_{\mathbb{C}}$, intertwining σ_0 with σ. (Equivalently, we require that $\phi_{\mathbb{R}}(\mathfrak{sl}(2,\mathbb{R})) \subset \mathfrak{g}$).*
3. *Conjugacy classes under K of Lie algebra homomorphisms $\phi_{\mathbb{R},\theta}$ from $\mathfrak{sl}(2,\mathbb{C})$ to $\mathfrak{g}_{\mathbb{C}}$, intertwining σ_0 with σ and θ_0 with θ.*
4. *Conjugacy classes under $K_{\mathbb{C}}$ of Lie algebra homomorphisms ϕ_θ from $\mathfrak{sl}(2,\mathbb{C})$ to $\mathfrak{g}_{\mathbb{C}}$, intertwining θ_0 with θ.*
5. *Orbits of $K_{\mathbb{C}}$ on the cone $\mathcal{N}_\theta$ of nilpotent elements in $\mathfrak{s}_{\mathbb{C}}$ (the -1 eigenspace of θ on $\mathfrak{g}_{\mathbb{C}}$).*

Here the bijections from (3) to (2) and (4) are given by the obvious inclusions; that from (2) to (1) sends $\phi_{\mathbb{R}}$ to $\phi_{\mathbb{R}}(E_0)$ (notation (4.1)(c)); and that from (4) to (5) sends ϕ_θ to $\phi_\theta(x_0)$.

This formulation incorporates ideas of Jacobson-Morozov, Kostant-Rallis, and Sekiguchi (among others); there is a more detailed discussion of the proof in [**VAV**], Theorem 6.4. The correspondence between (1) and (5) is the *Kostant-Sekiguchi correspondence* between real nilpotent orbits of G and nilpotent orbits of $K_{\mathbb{C}}$ on $\mathfrak{s}_{\mathbb{C}}$. If $\phi_{\mathbb{R},\theta}$ is a homomorphism as in (3), then the correspondence is

$$E = \phi_{\mathbb{R},\theta}(E_0) \leftrightarrow x = \phi_{\mathbb{R},\theta}(x_0)$$

(notation as in (4.1)(c,f)).

We now examine what Proposition 4.2 can tell us about the structure of isotropy groups and orbits.

Proposition 4.3. *Suppose ϕ is a Lie algebra homomorphism from $\mathfrak{sl}(2,\mathbb{C})$ to a complex reductive Lie algebra $\mathfrak{g}_{\mathbb{C}}$. Write*

$$E = \phi(E_0), \quad H = \phi(H_0), \quad F = \phi(F_0)$$

(notation as in (4.1)(c)).

1. *The operator $\mathrm{ad}(H)$ has integer eigenvalues; so if we write*
$$\mathfrak{g}_{\mathbb{C}}(p) = \{X \in \mathfrak{g}_{\mathbb{C}} \mid [H,X] = pX\},$$
then
$$\mathfrak{g}_{\mathbb{C}} = \sum_{p\in\mathbb{Z}} \mathfrak{g}_{\mathbb{C}}(p).$$
2. *Write*
$$\mathfrak{l} = \mathfrak{g}_{\mathbb{C}}(0), \qquad \mathfrak{u} = \sum_{p>0} \mathfrak{g}_{\mathbb{C}}(p).$$
Then $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is a Levi decomposition of a parabolic subalgebra of $\mathfrak{g}_{\mathbb{C}}$.
3. *The centralizer of E is graded by the decomposition in (1). More precisely,*
$$\mathfrak{g}_{\mathbb{C}}^E = \mathfrak{l}^E + \sum_{p>0} \mathfrak{g}_{\mathbb{C}}(p)^E = \mathfrak{l}^E + \mathfrak{u}^E.$$

4. *The subalgebra* $\mathfrak{l}^E = \mathfrak{g}_{\mathbb{C}}^{H,E}$ *is equal to* $\mathfrak{g}_{\mathbb{C}}^{\phi}$*, the centralizer in* $\mathfrak{g}_{\mathbb{C}}$ *of the image of* ϕ*. It is a reductive subalgebra of* $\mathfrak{g}_{\mathbb{C}}$*. Consequently the decomposition in (3) is a Levi decomposition of the algebraic Lie algebra* $\mathfrak{g}_{\mathbb{C}}^E$*.*

Parallel results hold if (E, H, F) *are replaced by*

$$x = \phi(x_0), \quad h = \phi(h_0), \quad y = \phi(y_0)$$

(notation as in (4.1)(f)).

Most of the assertions in this proposition follow from the representation theory of $\mathfrak{sl}(2)$, applied to the representation $\mathrm{ad} \circ \phi$ on $\mathfrak{g}_{\mathbb{C}}$. The rest are fairly easy; we omit the arguments.

Now we want analogous facts on the group level.

Proposition 4.4. *Suppose G is a real reductive Lie group, and $\phi_{\mathbb{R}}$ is a Lie algebra homomorphism from $\mathfrak{sl}(2,\mathbb{C})$ to $\mathfrak{g}_{\mathbb{C}}$, intertwining σ_0 with σ. Define E, H, and F in $\mathfrak{g}$ as in Proposition 4.3.*

1. *The elements E and F are nilpotent, and H is hyperbolic.*
2. *Define $L = G^H$ to be the isotropy group of the adjoint action at H, and $U = \exp(\mathfrak{u} \cap \mathfrak{g})$. Then $Q = LU$ is the parabolic subgroup of G associated to H (Proposition 2.15).*
3. *The stabilizer of E in the adjoint action is contained in Q, and respects the Levi decomposition:*
$$G^E = (L^E)(U^E).$$
4. *The subgroup $L^E = G^{H,E}$ is equal to $G^{\phi_{\mathbb{R}}}$, the centralizer in G of the image of $\phi_{\mathbb{R}}$. It is a reductive subgroup of G. The subgroup U^E is simply connected unipotent.*
5. *Suppose that $\phi_{\mathbb{R}} = \phi_{\mathbb{R},\theta}$ also intertwines θ_0 with θ. Then $G^{\phi_{\mathbb{R},\theta}}$ is preserved by θ, and we may take θ as a Cartan involution on this reductive group. In particular, G^E and $G^{\phi_{\mathbb{R},\theta}}$ have a common maximal compact subgroup*
$$K^{\phi_{\mathbb{R},\theta}} = (L \cap K)^E.$$

This proposition provides good information about the action of G on nilpotent elements in $\mathfrak{g}$. Eventually we will want the same information about the action of $K_{\mathbb{C}}$ on nilpotent elements in $\mathfrak{s}_{\mathbb{C}}$. Here it is.

Proposition 4.5. *Suppose G is a real reductive Lie group with Cartan decomposition $G = K\exp(\mathfrak{s})$. Write $K_{\mathbb{C}}$ for the complexification of K. Suppose that ϕ_θ is a Lie algebra homomorphism from $\mathfrak{sl}(2,\mathbb{C})$ to $\mathfrak{g}_{\mathbb{C}}$, intertwining θ_0 with θ. Define x, h, and y in $\mathfrak{g}$ as in Proposition 4.3.*

1. *The elements x and y are in $\mathfrak{s}_{\mathbb{C}}$ are nilpotent, $h \in \mathfrak{k}_{\mathbb{C}}$ is hyperbolic, and $ih \in \mathfrak{k}_{\mathbb{C}}$ is elliptic.*
2. *The parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ of $\mathfrak{g}$ constructed as in Proposition 4.3 using h, is preserved by θ.*
3. *Define $L_K = (K_{\mathbb{C}})^h$ to be the isotropy group of the adjoint action of $K_{\mathbb{C}}$ at h, and $U_K = \exp(\mathfrak{u} \cap \mathfrak{k}_{\mathbb{C}})$. Then $Q_K = L_K U_K$ is the parabolic subgroup of $K_{\mathbb{C}}$ associated to h (Proposition 2.15).*
4. *The stabilizer of x in the adjoint action of $K_{\mathbb{C}}$ is contained in Q_K, and respects the Levi decomposition:*
$$K_{\mathbb{C}}^x = (L_K^x)(U_K^x).$$

5. *The subgroup* $L_K^x = K_{\mathbb{C}}^{h,x}$ *is equal to* $K_{\mathbb{C}}^{\phi_\theta}$*, the centralizer in* $K_{\mathbb{C}}$ *of the image of* ϕ_θ*. It is a reductive algebraic subgroup of* $K_{\mathbb{C}}$*. The subgroup* U_K^x *is simply connected unipotent. In particular, the decomposition of (4) is a Levi decomposition of the algebraic group* $K_{\mathbb{C}}^x$*.*
6. *Suppose that* $\phi_\theta = \phi_{\mathbb{R},\theta}$ *also intertwines* σ_0 *with* σ*. Then* $K_{\mathbb{C}}^{\phi_{\mathbb{R},\theta}}$ *is preserved by* σ*, and we may take* σ *as complex conjugation for a compact real form of this reductive algebraic group. In particular,* $K_{\mathbb{C}}^x$ *and* $K_{\mathbb{C}}^{\phi_{\mathbb{R},\theta}}$ *have a common maximal compact subgroup*

$$K^{\phi_{\mathbb{R},\theta}} = L_K^x \cap K.$$

These are rather convoluted and technical statements, not easy to absorb at first. One way to think about them is this. The Cartan decomposition says that a reductive group G is topologically a product of a maximal compact subgroup K and a vector space. A consequence of this decomposition is that many nice homogeneous spaces G/H are isomorphic to vector bundles over $K/K \cap H$. Said a little more abstractly: a nice homogeneous space for G is a vector bundle over

$$(\text{maximal compact for } G)/(\text{maximal compact for isotropy group}).$$

In order to understand a homogeneous space topologically, we must therefore understand maximal compact subgroups of isotropy groups.

Proposition 4.4 computes a maximal compact subgroup of an isotropy group for a nilpotent adjoint orbit. The natural inference (which is actually true) is that this orbit $G \cdot E$ is isomorphic to a vector bundle over $K/K^{\phi_{\mathbb{R},\theta}}$; the isomorphism respects the action of K.

Proposition 4.5 analyzes $K_{\mathbb{C}} \cdot x \subset \mathfrak{s}_{\mathbb{C}}$. Notice that $K_{\mathbb{C}}$ has exactly the same maximal compact subgroup K as G; and the Proposition says that the maximal compact subgroups of isotropy groups also look the same. The inference (again true) is that $K_{\mathbb{C}} \cdot x$ is isomorphic to a vector bundle over $K/K^{\phi_{\mathbb{R},\theta}}$; the isomorphism again respects the action of K.

It is not impossibly difficult to calculate the vector bundles appearing in these two settings, and to see that they are isomorphic. The conclusion is that the orbits $G \cdot E$ and $K_{\mathbb{C}} \cdot x$ are diffeomorphic as manifolds with K action (assuming that they correspond under the Kostant-Sekiguchi correspondence). Vergne has given a more natural and direct proof of this diffeomorphism. Here is a statement.

Theorem 4.6 ([Vrg]). *Suppose* $G = K \cdot \exp(\mathfrak{s})$ *is a Cartan decomposition of a reductive group, and* $E \in \mathfrak{g}$ *and* $x \in \mathfrak{s}_{\mathbb{C}}$ *are nilpotent elements. Assume that the orbits* $G \cdot E$ *and* $K_{\mathbb{C}} \cdot x$ *correspond under the Kostant-Sekiguchi correspondence (Proposition 4.2). Then there is a* K*-equivariant diffeomorphism from* $G \cdot E$ *onto* $K_{\mathbb{C}} \cdot x$*.*

LECTURE 5
Quantizing the action of K

We would like to be able to attach to a nilpotent coadjoint orbit a unitary representation of G. The construction of parabolically induced representations (on hyperbolic orbits) suggests that we might look for a pre-Hilbert space endowed with a representation of G preserving the inner product. The precise construction used for cohomological induction (in the case of elliptic orbits) produces even less: a pre-Hilbert space endowed with a $(\mathfrak{g}, K)$-module structure. In the case of nilpotent orbits, we will be able to produce only a space with a representation of K (and even that only under an additional technical hypothesis). The action of $\mathfrak{g}$ must be added later (and we will usually not know how to do it).

So let us fix a nilpotent element

$$\lambda_n \in \mathfrak{g}^*, \tag{5.1)(a}$$

and write $E \in \mathfrak{g}$ for its image under the isomorphism of Proposition 2.7. According to Proposition 4.2, we can find a Lie algebra homomorphism $\phi_{\mathbb{R}}$ from $\mathfrak{sl}(2,\mathbb{R})$ to $\mathfrak{g}$ with $\phi_{\mathbb{R}}(E_0) = E$ (notation (4.1)(c)). After replacing λ_n by a conjugate under G, we may assume that

$$\phi_{\mathbb{R}} = \phi_{\mathbb{R},\theta} \text{ intertwines } \theta_0 \text{ and } \theta. \tag{5.1)(b}$$

Following the prescription of Proposition 4.2, we then define

$$x = \phi_{\mathbb{R},\theta}(x_0) \in \mathfrak{s}_{\mathbb{C}}. \tag{5.1)(c}$$

The isomorphism of Proposition 2.7 associates to x a linear functional

$$\lambda_\theta \in \mathfrak{s}_{\mathbb{C}}^*, \qquad \lambda_\theta(Y) = \langle x, Y \rangle \qquad (Y \in \mathfrak{g}). \tag{5.1)(d}$$

The element λ_θ is not uniquely determined by λ_n, but the orbit $K_{\mathbb{C}}\cdot\lambda_\theta$ is determined by $G\cdot\lambda_n$. Vergne's Theorem 4.6 provides a K-equivariant diffeomorphism

$$G\cdot\lambda_n \simeq K_{\mathbb{C}}\cdot\lambda_\theta. \tag{5.1)(e}$$

The discussion of elliptic orbits in section 3 showed that we should expect to need a little more than the orbit to construct a representation. Specifically, we should have also an "admissible orbit datum at λ_n" in the sense of Duflo. For the purposes of these lectures, defining admissible orbit data would be an unnecessary digression. What we need to understand is what becomes of admissible orbit data on $G\cdot\lambda_n$ after they are transferred to $K_{\mathbb{C}}\cdot\lambda_\theta$. Here is the key definition.

Definition 5.2. Suppose $\lambda_\theta \in \mathfrak{s}_{\mathbb{C}}^*$ is a nilpotent element. Write $K_{\mathbb{C}}^{\lambda_\theta}$ for the isotropy group at λ_θ. Define 2ρ to be the algebraic character of $K_{\mathbb{C}}^{\lambda_\theta}$ by which it acts on the top exterior power of the cotangent space at λ_θ to the orbit:

$$2\rho(k) = \det\left(\mathrm{Ad}^*(k)|_{(\mathfrak{k}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta})^*}\right) \qquad (k \in K_{\mathbb{C}}^{\lambda_\theta}).$$

The differential of 2ρ is a one-dimensional representation of the Lie algebra $\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta}$, which we denote also by $2\rho \in (\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta})^*$. We define $\rho \in (\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta})^*$ to be half of 2ρ. Explicitly,

$$\rho(Z) = \frac{1}{2}\operatorname{tr}(\mathrm{ad}^*(Z)|_{(\mathfrak{k}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta})^*}) \qquad (Z \in \mathfrak{k}_{\mathbb{C}}^{\lambda_\theta}).$$

An *admissible $K_{\mathbb{C}}$-orbit datum at λ_θ* is an irreducible algebraic representation (τ, V_τ) of $K_{\mathbb{C}}^{\lambda_\theta}$ whose differential is equal to ρ (times the identity on V_τ). The element λ_θ is called *admissible* if an admissible orbit datum exists.

Just as in the case of elliptic orbits, admissible orbit data may or may not exist. When they do exist, there is often a one-dimensional admissible datum (τ_0, V_{τ_0}). In that case it is easy to see that all admissible data are in one-to-one correspondence with irreducible representations of the group of connected components of $K_{\mathbb{C}}^{\lambda_\theta}$; the correspondence arises by tensoring with τ_0. If G is connected and simply connected, then this component group is just the fundamental group of the orbit $K_{\mathbb{C}} \cdot \lambda_\theta$.

Theorem 5.3 (Schwartz; see [**VAV**], Theorem 7.14)**.** *In the setting of (5.1), there is a natural bijection from Duflo's admissible orbit data at λ_n to admissible $K_{\mathbb{C}}$-orbit data at λ_θ. In particular, the orbit $G \cdot \lambda_n$ is admissible in Duflo's sense if and only if the orbit $K_{\mathbb{C}} \cdot \lambda_\theta$ is admissible.*

We now have in hand the ingredients necessary to imitate the construction of a quantization made for elliptic orbits: a complex structure on the orbit, and a holomorphic vector bundle. The difficulty is that the complex structure is preserved only by K, and not by the whole group G; so the space of holomorphic sections of the bundle is a representation only of K. Here is the definition. (The definition would still make sense without the codimension condition imposed here, but it is no longer a reasonable one.)

Definition 5.4. Suppose $\lambda_n \in \mathfrak{g}^*$ is a nilpotent element, and $\lambda_\theta \in \mathfrak{s}_{\mathbb{C}}^*$ corresponds to it under the Kostant-Sekiguchi correspondence. Fix an admissible orbit datum (τ, V_τ) at λ_θ, and let

$$\mathcal{V}_\tau = K_{\mathbb{C}} \times_{K_{\mathbb{C}}^{\lambda_\theta}} V_\tau$$

be the corresponding algebraic vector bundle on $K_{\mathbb{C}} \cdot \lambda_\theta$. Assume that the boundary of $\overline{K_{\mathbb{C}} \cdot \lambda_\theta}$ (that is, the complement of the open orbit) has complex codimension at least two. Define

$$X_K(\lambda_n, \tau) = \text{space of algebraic sections of } \mathcal{V}_\tau,$$

an algebraic representation of $K_{\mathbb{C}}$. We call this the *quantization of the K action on $G \cdot \lambda_n$* (for the admissible orbit datum τ).

What this definition amounts to is a desideratum for the quantization of the G action on $G \cdot \lambda_n$: that is, that whatever unitary representation $\pi_G(\lambda_n, \tau)$ we associate to these data, we should have

$$K\text{-finite part of } \pi_G(\lambda_n, \tau) \simeq X_K(\lambda_n, \tau). \tag{5.5}$$

One might justify this requirement just by analogy with the quantization of elliptic orbits in Proposition 3.5. One of the main points of [**VAV**] is to provide a stronger justification: roughly speaking, to show that any representation of G satisfying certain rather weak and natural requirements related to $G \cdot \lambda_n$ must come close to satisfying (5.5). (A more precise statement is at (12.2) in [**VAV**].) There is also a growing collection of examples that support (5.5); a recent one is in [**AHV**].

In the absence of the codimension condition in Definition 5.4, I do not know in general how to formulate a requirement like (5.5); that is, how to quantize the action of K. For the moment it is enough to notice that the condition is satisfied in a great many interesting examples (see Exercise 13). When G is complex, the coadjoint orbits are complex symplectic manifolds, and therefore of real dimension $4m$. Consequently the codimension condition is automatically satisfied in this case.

LECTURE 6
Associated graded modules

In section 5 we have attached to some admissible nilpotent coadjoint orbits certain candidates for the restrictions to K of corresponding unitary representations. In this section we will study the restrictions to K of arbitrary admissible representations, finding a description of more or less the same nature as the conjectural one from the orbit method. Our goal is to understand this general description so well that we can sometimes prove that it coincides with the conjectural one.

Definition 6.1. Suppose $G = K\exp(\mathfrak{s})$ is a Cartan decomposition of a reductive group. Write

$$\mathcal{M}(\mathfrak{g}, K) = \text{category of } (\mathfrak{g}, K)\text{-modules of finite length.}$$

The set of equivalence classes of irreducible $(\mathfrak{g}, K)$-modules is denoted $\widehat{G}$. Similarly, write

$$\mathcal{A}(K) = \text{category of admissible representations of } K.$$

These are the locally finite representations of K (every vector belongs to a finite-dimensional K-invariant subspace on which K acts continuously) with the property that each irreducible representation of K occurs finitely often. One of Harish-Chandra's basic theorems about $(\mathfrak{g}, K)$-modules says that any $(\mathfrak{g}, K)$-module of finite length belongs to $\mathcal{A}(K)$; that is, that there is a forgetful functor

$$\mathcal{M}(\mathfrak{g}, K) \to \mathcal{A}(K).$$

Recall that the *Grothendieck group* of an abelian category $\mathcal{A}$ is the abelian group $K\mathcal{A}$ generated by the objects of the category, modulo relations

$$[B] = [A] + [C]$$

whenever there is a short exact sequence

$$0 \to A \to B \to C \to 0.$$

Here and in general we write $[A] \in K\mathcal{A}$ for the class of the object A in the Grothendieck group.

Both of the module categories above are abelian, so they have Grothendieck groups. Every object in $\mathcal{M}(\mathfrak{g}, K)$ has finite length; it follows easily that $K\mathcal{M}(\mathfrak{g}, K)$

is a free abelian group with generators the irreducible $(\mathfrak{g}, K)$-modules. Here is another way to say the same thing: if M is any $(\mathfrak{g}, K)$-module of finite length, then

$$[M] = \sum_{\pi \in \widehat{G}} m_\pi(M)[V_\pi].$$

Here $m_\pi(M)$ is a non-negative integer, the number of times that π appears in a Jordan-Hölder series for M. This sum has only finitely many non-zero terms.

The situation for $\mathcal{A}(K)$ is a bit more subtle, since the objects do not have finite length. In this case the Grothendieck group is a direct *product* of copies of $\mathbb{Z}$, one for each irreducible representation of K. Explicitly, if N is any admissible representation of K, then

$$[N] = \sum_{\mu \in \widehat{K}} m_\mu(N)[V_\mu].$$

Here $m_\mu(N)$ is the multiplicity of μ in N, a non-negative integer; but the sum need not be finite.

Finally, write $K_{\mathbb{C}}$ for the complexification of K; this is a complex reductive algebraic group. The locally finite representations of K may be identified with the algebraic representations of $K_{\mathbb{C}}$. (Formally, the forgetful functor from the latter to the former is an equivalence of categories.) We may therefore replace K by $K_{\mathbb{C}}$ at will in these definitions without changing anything.

We are going to need one more abelian category and its Grothendieck group. We could introduce the category directly, but perhaps it is preferable to begin with the construction that motivates it. So suppose M is a $(\mathfrak{g}, K)$-module of finite length. An easy induction on the length shows that M is generated as a $(\mathfrak{g}, K)$-module by some finite-dimensional subspace S; in fact we can choose S to be of dimension at most the length of M. Next, the definition of local finiteness for the action of K ensures that S is contained in a finite-dimensional K-invariant subspace M_0 of M. The compatibility of the K and $\mathfrak{g}$ actions ensures that $U(\mathfrak{g})M_0$ is K-invariant; since it is also $\mathfrak{g}$ invariant, it must be a $(\mathfrak{g}, K)$-submodule of M. Since it contains S, it must be all of M:

$$M = U(\mathfrak{g})M_0. \tag{6.2)(a}$$

Now the enveloping algebra $U(\mathfrak{g})$ has an increasing filtration

$$\mathbb{C} = U_0(\mathfrak{g}) \subset \mathbb{C} + \mathfrak{g}_{\mathbb{C}} = U_1(\mathfrak{g}) \subset U_2(\mathfrak{g}) \subset \cdots; \tag{6.2)(b}$$

here $U_n(\mathfrak{g})$ is the span of products of at most n elements of $\mathfrak{g}$. The subspaces $U_n(\mathfrak{g})$ are finite-dimensional and $\mathrm{Ad}(K)$-invariant. Now define

$$M_n = U_n(\mathfrak{g})M_0. \tag{6.2)(c}$$

This is a finite-dimensional K-invariant subspace of M. Because $U_p(\mathfrak{g})U_q(\mathfrak{g}) \subset U_{p+q}(\mathfrak{g})$, we have

$$U_p(\mathfrak{g})M_q \subset M_{p+q}. \tag{6.2)(d}$$

The most important fact about the universal enveloping algebra is the Poincaré-Birkhoff-Witt theorem. There are many equivalent formulations (often involving a basis for $U(\mathfrak{g})$) but for us the following coordinate-free version is best.

Theorem 6.3 (**[Hum]**, Theorem 17.3)**.** *If $U(\mathfrak{g})$ is given the structure of filtered algebra described above, then the associated graded algebra is naturally isomorphic to $S(\mathfrak{g})$, the symmetric algebra on the complex vector space $\mathfrak{g}_{\mathbb{C}}$. We write*

$$\sigma_n: U_n(\mathfrak{g})/U_{n-1}(\mathfrak{g}) \to S^n(\mathfrak{g})$$

for the isomorphism; here S^n denotes homogeneous polynomials of degree n.

The notation is chosen because σ_n behaves like a symbol map for differential operators. (This is no accident: if $U(\mathfrak{g})$ is identified with left-invariant differential operators on G, then the filtration we have is precisely the degree filtration for operators, and σ can be identified with the symbol of the differential operator.) One can interpret the theorem as saying that multiplication in $U(\mathfrak{g})$ is like multiplication of polynomials, up to lower order terms.

Now we want to understand not $U(\mathfrak{g})$ itself, but the $U(\mathfrak{g})$-module M. We will see that the module structure is like that of a module for the polynomial ring $S(\mathfrak{g})$, modulo lower order terms. Here is the result.

Lemma 6.4. *Suppose M is a $(\mathfrak{g}, K)$-module of finite length. Construct a filtration of M as in (6.2) above. Write $\operatorname{gr} M$ for the associated graded vector space:*

$$(\operatorname{gr} M)_n = M_n/M_{n-1} \qquad (n \geq 0);$$

here we define $M_{-1} = 0$. Each of these spaces inherits from M a representation of the group K; so $\operatorname{gr} M$ is an admissible locally finite representation of K. Because of (6.2)(d), there are well-defined maps

$$(\operatorname{gr} U)_p \otimes (\operatorname{gr} M)_q \to (\operatorname{gr} M)_{p+q},$$

$$(u + U_{p-1}) \otimes (m + M_{q-1}) \mapsto um + M_{p+q-1} \qquad (u \in U_p(\mathfrak{g}), m \in M_q).$$

In light of Theorem 6.3, these may be interpreted as maps

$$S^p(\mathfrak{g}) \otimes (\operatorname{gr} M)_q \to (\operatorname{gr} M)_{p+q}.$$

These maps make $\operatorname{gr} M$ into a finitely generated graded module for $S(\mathfrak{g})$, generated by the degree 0 subspace M_0. The action of $S(\mathfrak{g})$ and the representation of K on $\operatorname{gr} M$ satisfy the compatibility condition

$$k \cdot (p \cdot m) = (\operatorname{Ad}(k)p) \cdot (k \cdot m) \qquad (k \in K, p \in S(\mathfrak{g}), m \in \operatorname{gr} M).$$

This lemma is easier to prove than it is to state; so having assumed the burden of stating it, we may safely leave the proof to the reader.

At this point we begin to take advantage of the geometric aspects of algebraic geometry. Recall first of all that the maximal ideals in the polynomial ring $S(\mathfrak{g})$ correspond precisely to the complex-linear functionals on $\mathfrak{g}$. We will write a little informally

$$\operatorname{Spec} S(\mathfrak{g}) = \mathfrak{g}^*_{\mathbb{C}}, \tag{6.5)(a}$$

having in mind the identification of prime ideals with irreducible subvarieties of $\mathfrak{g}^*_{\mathbb{C}}$. If X is any subset of $S(\mathfrak{g})$, we write

$$\mathcal{V}(X) = \{\lambda \in \mathfrak{g}^*_{\mathbb{C}} \mid p(\lambda) = 0, \text{all } p \in X\}, \tag{6.5)(b}$$

the *associated variety of* X. This is the same as the associated variety of the ideal generated by X. If N is a finitely generated $S(\mathfrak{g})$-module, we can define

$$\mathcal{V}(N) = \mathcal{V}(\operatorname{Ann} N), \tag{6.5)(c}$$

the *associated variety of* N. A standard argument from commutative algebra shows that $\mathcal{V}(N)$ coincides with the support of N, the set of all prime ideals for which the localization N_P is non-zero. (For details we refer to [**VAV**], section 2.) Our next goal is to say something about the associated variety of $\operatorname{gr} M$ in the setting of Lemma 6.4. By definition, that means we need to find some elements in $\operatorname{Ann} \operatorname{gr} M$.

Because K preserves M_n, the differentiated action of $\mathfrak{k}$ also preserves M_n. By the definition of $(\mathfrak{g}, K)$-module, this is just the restriction to $\mathfrak{k}$ of the $\mathfrak{g}$ action. By the definition of the $S(\mathfrak{g})$ action in Lemma 6.4, it follows that $\mathfrak{k}$ acts by 0 on $\operatorname{gr} M$. (Elements of $\mathfrak{g}$ are supposed to raise degree by 1 in M.) Consequently

$$\operatorname{gr} M \text{ is a finitely generated graded } S(\mathfrak{g}/\mathfrak{k}) \text{ module.} \tag{6.6)(a}$$

In terms of associated varieties, this says

$$\mathcal{V}(\operatorname{gr} M) \subset (\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}})^* \subset \mathfrak{g}_{\mathbb{C}}^*. \tag{6.6)(b}$$

Write $\mathfrak{Z}(\mathfrak{g})$ for the center of the universal enveloping algebra $U(\mathfrak{g})$. This is a polynomial algebra on $\operatorname{rank} \mathfrak{g}$ generators. It inherits from $U(\mathfrak{g})$ a filtration; the associated graded algebra is

$$\operatorname{gr} \mathfrak{Z}(\mathfrak{g}) \simeq S(\mathfrak{g})^{G_0}, \tag{6.6)(c}$$

the algebra of polynomials invariant under the adjoint action of G_0. Because M is a $U(\mathfrak{g})$-module of finite length, it must be annihilated by an ideal $I \subset \mathfrak{Z}(\mathfrak{g})$ of finite codimension. (The codimension of I may be bounded by the length of M as a $U(\mathfrak{g})$-module.) Since $I \cdot M = 0$, it follows at once that

$$(\operatorname{gr} I) \cdot (\operatorname{gr} M) = 0. \tag{6.6)(d}$$

Now $\operatorname{gr} I$ is a graded ideal of finite codimension in $S(\mathfrak{g})^{G_0}$. By an elementary argument, $\operatorname{gr} I$ must contain some power of the augmentation ideal $J = \sum_{n>0} S^n(\mathfrak{g})^{G_0}$:

$$\operatorname{gr} I \supset J^N, \qquad \text{some } N > 0. \tag{6.6)(e}$$

In terms of associated varieties, this says

$$\mathcal{V}(\operatorname{gr} M) \subset \mathcal{V}(J^N) = \mathcal{V}(J). \tag{6.6)(f}$$

Now the associated variety of J is not so obvious as that of $\mathfrak{k}$. Fortunately Kostant has computed it; here is his result.

Theorem 6.7 ([**VAV**], Theorem 5.7)**.** *Suppose G is a reductive Lie group, and $J \subset S(\mathfrak{g})$ is the collection of G_0-invariant polynomials without constant term. Then the associated variety of J is the cone $\mathcal{N}^*$ of nilpotent elements in $\mathfrak{g}_{\mathbb{C}}^*$.*

Suppose in particular that M is a $(\mathfrak{g}, K)$-module of finite length, endowed with a filtration as in (6.2). Then the associated variety of $\operatorname{gr} M$ *is contained in the cone $\mathcal{N}_\theta^*$ of nilpotent elements in $(\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}})^*$.*

Now we can define the new abelian category mentioned after Definition 6.1.

Definition 6.8. Let $\mathcal{C}(\mathfrak{g}, K)$ be the category of finitely generated $S(\mathfrak{g}/\mathfrak{k})$-modules N carrying locally finite representations of K, and subject to the following requirements: first, the compatibility requirement

$$k \cdot (p \cdot n) = (\mathrm{Ad}(k)p) \cdot (k \cdot m) \qquad (k \in K, p \in S(\mathfrak{g}), n \in N);$$

and second, the support requirement

$$\mathcal{V}(N) \subset \mathcal{N}_\theta^*$$

(the cone of nilpotent elements in $(\mathfrak{g}_\mathbb{C}/\mathfrak{k}_\mathbb{C})^*$). This is an abelian category. In light of Kostant's Theorem 6.7, the support requirement is equivalent to requiring J (the set of G_0-invariant polynomials without constant term) to act nilpotently on N.

Because of the remarks at the end of Definition 6.1, it is equivalent to consider the corresponding category $\mathcal{C}(\mathfrak{g}, K_\mathbb{C})$ of modules with an algebraic action of $K_\mathbb{C}$.

Proposition 6.9. *Suppose N is an object in $\mathcal{C}(\mathfrak{g}, K)$. Then N is admissible (finite multiplicities) as a representation of K. In particular, there is a forgetful functor*

$$\mathcal{C}(\mathfrak{g}, K) \to \mathcal{A}(K).$$

Sketch of proof. An easy argument reduces the lemma to the case G connected and

$$N = S(\mathfrak{g})/(\mathfrak{k} + J)S(\mathfrak{g}).$$

This case was treated by Kostant and Rallis in [**KR**], who showed that the representation μ of K appears with multiplicity equal to the dimension of the space of M-fixed vectors in μ. (Here $M \subset K$ comes from the Langlands decomposition MAN of a minimal parabolic subgroup of G.) □

We would like the operation "associated graded" to define something like a functor from $\mathcal{M}(\mathfrak{g}, K)$ (Definition 6.1) to $\mathcal{C}(\mathfrak{g}, K)$. The difficulty is that our definition of $\mathrm{gr}\, M$ required the construction of a filtration of M, and the construction required at least one non-functorial choice (of the finite-dimensional generating subspace S of M). Fortunately this choice does not affect $\mathrm{gr}\, M$ too much.

Proposition 6.10 ([**VAV**], Proposition 2.2)**.** *Suppose M is a $(\mathfrak{g}, K)$-module of finite length, and $\{M_n\}$ and $\{M'_n\}$ are two filtrations constructed as in (6.2). Then the corresponding associated graded modules $\mathrm{gr}\, M$ and $\mathrm{gr}'\, M$ belong to $\mathcal{C}(\mathfrak{g}, K)$. They define the same class in the Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)$:*

$$[\mathrm{gr}\, M] = [\mathrm{gr}'\, M].$$

More precisely, each associated graded module has a finite filtration so that the corresponding subquotients (after rearrangement) are pairwise isomorphic in $\mathcal{C}(\mathfrak{g}, K)$. Consequently there is a well-defined group homomorphism

$$[gr]: K\mathcal{M}(\mathfrak{g}, K) \to K\mathcal{C}(\mathfrak{g}, K).$$

The corresponding forgetful functors to $K\mathcal{A}(K)$ make a commutative triangle. That is, the K-multiplicities in M and $\mathrm{gr}\, M$ coincide.

That $\mathrm{gr}\, M$ belongs to $\mathcal{C}(\mathfrak{g}, K)$ follows from Lemma 6.4, Theorem 6.7, and (6.6)(a). The last statement about K-multiplicities is easy. For the rest of the proof, we refer to [**VAV**].

LECTURE 7
A good basis for associated graded modules

Proposition 6.10 provides a very useful way of working with the Grothendieck group of finite-length Harish-Chandra modules. Here is why.

Theorem 7.1. *The homomorphism* [gr] *of Proposition 6.10 is a surjection from the Grothendieck group of Harish-Chandra modules onto the Grothendieck group of finitely generated compatible* $(S(\mathfrak{g}/\mathfrak{k}), K)$*-modules supported on the nilpotent cone. Its kernel is precisely the group of virtual Harish-Chandra modules with* K*-multiplicities equal to zero; that is, the kernel of the forgetful functor*

$$K\mathcal{M}(\mathfrak{g}, K) \to K\mathcal{A}(K).$$

Corollary 7.2. *The possible* K*-multiplicities for virtual Harish-Chandra modules (that is, the image of the forgetful functor in Theorem 7.1) are precisely the* K*-multiplicities in virtual* $(S(\mathfrak{g}/\mathfrak{k}), K)$*-modules supported on the nilpotent cone.*

We are not going to give a detailed proof of Theorem 7.1. For the moment we are concerned mostly with Corollary 7.2. In order to extract interesting information about Harish-Chandra modules from this Corollary, we need to understand better the Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)$. This is a subtle matter, because the objects of $\mathcal{C}(\mathfrak{g}, K)$ (unless they are finite-dimensional) do not have finite length. (The Noetherian property guarantees that ascending chains of submodules must be finite, but infinite descending chains are easy to construct.) We cannot hope to use irreducible objects as a basis for the Grothendieck group. Instead we will begin with a theorem of Kostant and Rallis.

Theorem 7.3 (Kostant-Rallis; see [**VAV**], Corollary 5.22)**.** *Suppose* G *is a real reductive Lie group with maximal compact subgroup* K*, and* $K_{\mathbb{C}}$ *is the complexification of* K*. Write* $\mathcal{N}_\theta^*$ *for the cone of nilpotent elements in* $(\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}})^*$*. Then* $K_{\mathbb{C}}$ *acts on* $\mathcal{N}_\theta^*$ *with finitely many orbits.*

Choose representatives $\lambda_1, \ldots, \lambda_r$ for these orbits, and define

$$H_i = K_{\mathbb{C}}^{\lambda_i} \tag{7.4)(a}$$

to be the isotropy group at λ_i. We would like to disassemble the category $\mathcal{C}(\mathfrak{g}, K)$ using these orbits. If the orbits were closed, there would be no difficulty: the category would be a direct sum of subcategories consisting of modules supported on a single orbit. Each of these categories in turn would be equivalent to the category of finite-dimensional representations of an isotropy group H_i.

However, the orbits are not closed, and no such direct sum decomposition is possible. What we can do instead is filter the category $\mathcal{C}(\mathfrak{g}, K)$ using supports of modules. If an orbit $K_{\mathbb{C}} \cdot \lambda_j$ appears in the support of a module, and λ_i is in the closure of $K_{\mathbb{C}} \cdot \lambda_j$, then λ_i must also appear in the support. The support of any non-zero module N in $\mathcal{C}(\mathfrak{g}, K)$ may therefore be written uniquely as a union of orbit closures $\overline{K_{\mathbb{C}} \cdot \lambda_i}$, where λ_i is not in the closure of any other orbit in the support. These orbit closures are the components of $\mathcal{V}(N)$ as an algebraic variety.

This discussion is intended to motivate the following definition. For each i, define a subcategory of $\mathcal{C}(\mathfrak{g}, K)$ as follows.

$$\mathcal{C}(\mathfrak{g}, K)_i = \{N \in \mathcal{C}(\mathfrak{g}, K) \mid \lambda_i \in \left(\overline{K_{\mathbb{C}} \cdot \lambda_j} - K_{\mathbb{C}} \cdot \lambda_j\right) \Rightarrow \lambda_j \notin \mathcal{V}(N)\}. \qquad (7.4)(b)$$

These are the modules for which $\overline{K_{\mathbb{C}} \cdot \lambda_i}$ is a component of $\mathcal{V}(N)$, possibly with multiplicity zero. Next

$$\mathcal{C}(\mathfrak{g}, K)_i^0 = \{N \in \mathcal{C}(\mathfrak{g}, K) \mid \lambda_i \notin \mathcal{V}(N)\}. \qquad (7.4)(c)$$

These are the modules (automatically in $\mathcal{C}(\mathfrak{g}, K)_i$) for which $K_{\mathbb{C}} \cdot \lambda_i$ does not meet the associated variety.

Now any finitely generated module (say for a finitely generated commutative algebra over an algebraically closed field) is something like a vector bundle over a dense open set in its support. We should therefore expect modules in $\mathcal{C}(\mathfrak{g}, K)_i$ to look like $K_{\mathbb{C}}$-equivariant vector bundles on $K_{\mathbb{C}} \cdot \lambda_i \simeq K_{\mathbb{C}}/H_i$. Such a vector bundle is the same thing as an algebraic representation of H_i. The conclusion is that modules in $\mathcal{C}(\mathfrak{g}, K)_i$ should correspond to representations of H_i, with modules in $\mathcal{C}(\mathfrak{g}, K)_i^0$ corresponding to the zero representation. Making these ideas precise requires some passage to associated graded objects. (For example, a module is not necessarily annihilated by the ideal of its support, but only by some power of that ideal. We need a finite filtration of the module so that the subquotients are actually annihilated by the ideal of the support.) As a result, we get not a representation of H_i but a rather a class in some Grothendieck group. Here is some formalism useful for making this precise.

Lemma 7.5. *Suppose H is a complex algebraic group, with Levi decomposition $H = LU$. (Here U is the unipotent radical of H, a connected normal subgroup, and L is reductive.) Write $\mathcal{F}(H)$ for the category of finite-dimensional algebraic representations of H, and $\widehat{H}$ for the set of equivalence classes of irreducible algebraic representations of H.*

1. *Every irreducible representation of H is trivial on U; so restriction to L defines a bijection $\widehat{H} \simeq \widehat{L}$.*
2. *The irreducible representations of H constitute a basis for the Grothendieck group $K\mathcal{F}(H)$. In particular, the functor of restriction to L defines an isomorphism $K\mathcal{F}(H) \simeq K\mathcal{F}(L)$.*

The class in $K\mathcal{F}(H)$ of a representation of H is called a *genuine virtual representation.* According to the lemma, this is the same thing as a formal finite sum of irreducible representations.

Proposition 7.6 ([VAV], Theorem 2.13). *In the setting of (7.4), attached to any module $N \in \mathcal{C}(\mathfrak{g}, K)_i$ there is a genuine virtual representation $\chi(\lambda_i, N) \in K\mathcal{F}(H_i)$. This correspondence descends to an isomorphism of Grothendieck groups*

$$K\mathcal{C}(\mathfrak{g}, K)_i / K\mathcal{C}(\mathfrak{g}, K)_i^0 \simeq K\mathcal{F}(H_i).$$

There are various ways to build this into a statement about the full Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)$. One simple one is to define

$$\mathcal{C}(\mathfrak{g}, K)_{[m]} = \{N \in \mathcal{C}(\mathfrak{g}, K) \mid \dim \mathcal{V}(N) \leq m\}. \tag{7.7}$$

This is a finite increasing family of subcategories of $\mathcal{C}(\mathfrak{g}, K)$; if we define the empty set to have dimension -1, then $\mathcal{C}(\mathfrak{g}, K)_{[-1]}$ is zero. These subcategories induce a finite increasing filtration of the Grothendieck group, and we can describe the subquotients.

Corollary 7.8. *If $\dim K_{\mathbb{C}} \cdot \lambda_i = m$, then $\mathcal{C}(\mathfrak{g}, K)_{[m]} \subset \mathcal{C}(\mathfrak{g}, K)_i$. The maps of Proposition 7.6 induce isomorphisms of Grothendieck groups*

$$K\mathcal{C}(\mathfrak{g}, K)_{[m]} / K\mathcal{C}(\mathfrak{g}, K)_{[m-1]} \simeq \sum_{\substack{i \\ \dim K_{\mathbb{C}} \cdot \lambda_i = m}} K\mathcal{F}(H_i).$$

The Grothendieck groups $K\mathcal{F}(H_i)$ all have natural bases, parametrized by irreducible representations of H_i. We would like to lift these back to a basis for the Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)_i / K\mathcal{C}(\mathfrak{g}, K)_i^0$. The following proposition says that this is always possible, and sometimes can be done canonically.

Proposition 7.9. *In the setting of (7.4), suppose that (τ, V_τ) is an irreducible representation of the stabilizer H_i of λ_i. Then there is an object $N(\lambda_i, \tau) \in \mathcal{C}(\mathfrak{g}, K)$ with the following properties. First, $\mathcal{V}(N(\lambda_i, \tau)) = \overline{K_{\mathbb{C}} \cdot \lambda_i}$. In particular, $N(\lambda_i, \tau)$ belongs to $\mathcal{C}(\mathfrak{g}, K)_i$. Second, the map of Proposition 7.6 sends $N(\lambda_i, \tau)$ to $\chi(\lambda_i, N(\lambda_i, \tau)) = [\tau]$.*

If we choose such an object $N(\lambda_i, \tau)$ for every irreducible representation of every H_i, then the classes $[N(\lambda_i, \tau)]$ constitute a basis for the Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)$.

Suppose that the boundary of $\overline{K_{\mathbb{C}} \cdot \lambda_i}$ (that is, the complement of the orbit in its closure) has codimension at least two. Then there is a canonical choice for $N(\lambda_i, \tau)$. As representations of $K_{\mathbb{C}}$, we have

$$N(\lambda_i, \tau) \simeq \operatorname{Ind}_{H_i}^{K_{\mathbb{C}}}(\tau)$$

in this case.

Proof. We begin with the algebraic vector bundle

$$\mathcal{V}_\tau = K_{\mathbb{C}} \times_{H_i} V_\tau \tag{7.10(a)}$$

over the orbit $K_{\mathbb{C}} \cdot \lambda_i \simeq K_{\mathbb{C}}/H_i$. Define

$$M(\lambda_i, \tau) = \text{space of algebraic sections of } \mathcal{V}_\tau. \tag{7.10(b)}$$

This space carries an algebraic representation of $K_{\mathbb{C}}$; equivalently, a locally finite representation of K. At the same time, it is a module for $S(\mathfrak{g}/\mathfrak{k})$: a polynomial p acts by multiplication by the restriction of p to the orbit. (This statement is supposed to be elementary and clear, and from a geometric point of view it should be; but for a more group-oriented person there is a possibility of confusion. Such a person may prefer to identify a section m of $\mathcal{V}_\tau$ with an algebraic function

$$\mu: K_{\mathbb{C}} \to V_\tau, \qquad \mu(kh) = \tau(h^{-1})\mu(k) \qquad (k \in K_{\mathbb{C}}, h \in K_{\mathbb{C}}^{\lambda_i}). \tag{7.10)(c}$$

How does p act on such a function? To see that, we define from p a complex-valued function

$$\pi: K_{\mathbb{C}} \to \mathbb{C}, \qquad \pi(k) = p(k \cdot \lambda_i). \tag{7.10)(d}$$

The function π is right-invariant under $K_{\mathbb{C}}^{\lambda_i}$. Consequently $\pi\mu$ is a function from $K_{\mathbb{C}}$ to V_τ satisfying the same transformation law as μ; that is, $\pi\mu$ may be regarded as a section of $\mathcal{V}_\tau$. This is the section $p \cdot m$.)

The compatibility condition of Definition 6.8 is very easy to check. Because the orbit is quasiaffine (open in an affine algebraic variety) the bundle has many sections. In particular, the sections span the fiber at each point:

$$\{m(\lambda_i) \mid m \in M(\lambda_i, \tau)\} = V_\tau. \tag{7.10)(e}$$

(In the more group-theoretic description (7.10)(c) of sections of $\mathcal{V}_\tau$, the left side here is

$$\{\mu(e) \mid \mu: K_{\mathbb{C}} \to V_\tau \text{ as in } (7.10)(c)\}.$$

This description does not make it clear that we get all of V_τ from algebraic functions μ; for that the algebraic geometry point of view is essential.)

The only thing wrong with $M(\lambda_i, \tau)$ is that it need not be finitely generated as a module for $S(\mathfrak{g}/\mathfrak{k})$. To get around this, we choose a finite-dimensional K-invariant subspace M_0 of $M(\lambda_i, \tau)$ with the property that

$$\{m(\lambda_i) \mid m \in M_0\} = V_\tau. \tag{7.10)(f}$$

Such a subspace exists by (7.10)(e). Define

$$N(\lambda_i, \tau) = S(\mathfrak{g}/\mathfrak{k}) \cdot M_0, \tag{7.10)(g}$$

a finitely generated K-invariant submodule of $M(\lambda_i, \tau)$. It is not difficult to show that $N(\lambda_i, \tau)$ satisfies the requirements of the proposition. If we identify $N(\lambda_i, \tau)$ with a coherent sheaf of modules over $\overline{K_{\mathbb{C}} \cdot \lambda_i}$, then $M(\lambda_i, \tau)$ may be identified with the sections over the open set $K_{\mathbb{C}} \cdot \lambda_i$:

$$M(\lambda_i, \tau) \simeq N(\lambda_i, \tau)(K_{\mathbb{C}} \cdot \lambda_i). \tag{7.10)(h}$$

This description and some commutative algebra (see [**EGA**], Proposition 5.11.1) imply that $M(\lambda_i, \tau)$ is finitely generated if and only if the boundary of $\overline{K_{\mathbb{C}} \cdot \lambda_i}$ has codimension at least 2. When that condition is satisfied, we may therefore take $N(\lambda_i, \tau) = M(\lambda_i, \tau)$. This choice is in a certain sense maximal, if we require also that $N(\lambda_i, \tau)$ have no embedded associated primes; that is, no submodules with much larger annihilators. $\square$

Exercise 14 looks at the simplest example of this construction.

LECTURE 8
Proving unitarity

In section 7, we saw how to use the geometry of the nilpotent cone $\mathcal{N}_\theta^*$ to control the possible K-types of Harish-Chandra modules. In fact the work of Harish-Chandra and Langlands already provided a large and powerful array of tools for controlling the possible K-types of Harish-Chandra modules. In this section we will describe some of that classical theory, and see how comparing the two approaches can lead to constructions of unitary representations.

To begin, we need the notion of representations with real infinitesimal character. A precise definition appears in Definition 8.5 below. For most purposes it is enough to understand some properties of real infinitesimal character. First, discrete series representations have real infinitesimal character. Second, suppose that $P = MAN$ is a parabolic subgroup of G, δ is a representation of M with real infinitesimal character, and ν is a character of A. Then $\mathrm{Ind}_P^G(\delta \otimes \nu \otimes 1)$ has real infinitesimal character if and only if ν is real-valued.

Theorem 8.1. *Suppose G is a reductive Lie group, and M is an irreducible tempered Harish-Chandra module with real infinitesimal character. Then M has a unique lowest K-type $\mu(M)$, which occurs with multiplicity one. This defines a bijection from the set of equivalence classes of irreducible tempered Harish-Chandra modules with real infinitesimal character onto $\widehat{K}$. We may write $M(\mu)$ for the representation corresponding to $\mu \in \widehat{K}$.*

For linear groups in Harish-Chandra's class, this theorem is essentially proved in [**VG**]. The extension to general reductive groups is straightforward; what is perhaps surprising for experts is that nothing goes wrong.

Theorem 8.2. *The classes of irreducible tempered representations with real infinitesimal character form a basis for the quotient group*

$$K\mathcal{M}(\mathfrak{g}, K)/(\textit{kernel of restriction to } K).$$

By Theorem 7.1, the images

$$\{[\operatorname{gr} M] \mid M \textit{ tempered irreducible with real infinitesimal character}\}$$

form a basis of $K\mathcal{C}(\mathfrak{g}, K)$.

Even in the simplest cases, the bases provided by Proposition 7.9 and Theorem 8.2 are completely different. Exercise 15 looks at an example.

We want to use properties of the change-of-basis matrix from the basis of Proposition 7.9 to the one of Theorem 8.2 (for the Grothendieck group $K\mathcal{C}(\mathfrak{g}, K)$) in order to make deductions about unitary representations. (Of course it is not yet clear how we can hope to do that.) Because one basis is parametrized by irreducible representations of K, it is convenient to make use of highest weight theory for K. We therefore fix a Cartan subalgebra

$$\mathfrak{t} \subset \mathfrak{k}, \tag{8.3)(a}$$

and define

$$T = \{t \in K \mid \mathrm{Ad}(t)|_{\mathfrak{t}} = 1\}, \tag{8.3)(b}$$

$$T' = \{n \in K \mid \mathrm{Ad}(n)(\mathfrak{t}) = \mathfrak{t}\} = N_K(T), \tag{8.3)(c}$$

$$W(K,T) = T'/T. \tag{8.3)(d}$$

We may regard this Weyl group as a group of automorphisms of $\mathfrak{t}$ or of T. The identity component of T is a compact torus, but T may be non-abelian. We can also construct from $\mathfrak{t}$ a fundamental Cartan subalgebra and subgroup of G:

$$\mathfrak{h}^f = \{Z \in \mathfrak{g} \mid [Z, \mathfrak{t}] = 0\}, \qquad H^f = \{h \in G \mid \mathrm{Ad}(h)|_{\mathfrak{h}^f} = 1\}. \tag{8.3)(e}$$

Proposition 8.4. *Suppose $\mu \in \widehat{K}$. The restriction of μ to T (cf. (8.3)) contains a unique $W(K,T)$ orbit of extremal representations. If $\mu_0 \in \widehat{T}$ is such an extremal representation, write $\mu_1 \in i\mathfrak{t}^*$ for its differential. These weights form a single $W(K,T)$ orbit in $i\mathfrak{t}^*$. The corresponding set of weights $\lambda(\mu_1) \in i\mathfrak{t}^*$ constructed in [**VG**] is therefore a single $W(K,T)$ orbit. The infinitesimal character of $M(\mu)$ is represented by the weights in this orbit, extended to be θ-fixed complex linear functionals on $\mathfrak{h}^f$.*

We will not recall here the construction of $\lambda(\mu_1)$ in [**VG**]. One case is very simple, however: if G is a complex reductive group, then $\lambda(\mu_1) = \mu_1$.

We will be comparing "lengths" of infinitesimal characters, using the corresponding weights in $(\mathfrak{h}^f)^*$. A minor technical difficulty is that the bilinear form of Proposition 2.7 is not (after complexification) positive definite on $(\mathfrak{h}^f)^*$. The form *is* positive definite on the real span of the roots; so it is convenient to use something like that restriction to measure length. Here is the relevant definition, taken from [**VG**], Definition 5.4.11.

Definition 8.5. Suppose $\mathfrak{h}^f$ is a fundamental Cartan subalgebra as in (8.3). Write

$$\mathfrak{t} = \mathfrak{h}^f \cap \mathfrak{k}, \qquad \mathfrak{a}^f = \mathfrak{h}^f \cap \mathfrak{s}$$

for its decomposition into the $+1$ and -1 eigenspaces of the Cartan involution θ. The *canonical real part of* $\mathfrak{h}^f_{\mathbb{C}}$ is the subspace

$$\mathrm{RE}\, \mathfrak{h}^f_{\mathbb{C}} = i\mathfrak{t} + \mathfrak{a}^f.$$

(Notice that the roots of $\mathfrak{h}^f_{\mathbb{C}}$ take real values on this subspace, and that the form of Proposition 2.7 is real-valued and positive definite there.) The *canonical real part of* $\xi \in (\mathfrak{h}^f_{\mathbb{C}})^*$ is the element $\mathrm{RE}\,\xi$ characterized by the requirement that $\mathrm{RE}\,\xi$ takes real values on $\mathrm{RE}\,\mathfrak{h}^f_{\mathbb{C}}$, and $\xi - \mathrm{RE}\,\xi$ takes purely imaginary values there.

An infinitesimal character is called *real* if it is represented by a weight ξ equal to its canonical real part; that is, if ξ takes real values on $\mathrm{RE}\,\mathfrak{h}^f_{\mathbb{C}}$.

Definition 8.6. Suppose that $[N] \in K\mathcal{C}(\mathfrak{g}, K)$. Rewrite $[N]$ in the basis of Theorem 8.2 as

$$[N] = \sum_{\mu \in \widehat{K}} m_\mu(N)[M(\mu)].$$

Attach to $[N]$ a finite collection of $W(K,T)$ orbits in $i\mathfrak{t}^*$ by

$$\Lambda([N]) = \{\lambda(\mu) \mid \mu \in \widehat{K}, m_\mu(N) \neq 0\}.$$

(Recall that $\lambda(\mu)$ represents the infinitesimal character of $M(\mu)$.) *This collection is non-empty if and only if* $[N] \neq 0$. If $[N] \neq 0$, we can therefore define

$$\| [N] \| = \max_{\lambda \in \Lambda([N])} \|\lambda\|,$$

the *infinitesimal character size of* $[N]$.

The definition is interesting (and the terminology is justified) because of the following result.

Proposition 8.7. *Suppose* $[M] \in K\mathcal{M}(\mathfrak{g}, K)$ *is a virtual Harish-Chandra module such that* $[\operatorname{gr} M] = [N]$ *is a non-zero element of* $K\mathcal{C}(\mathfrak{g}, K)$. *Write*

$$[M] = \sum_{\pi \in \widehat{G}} m_\pi(M)[V_\pi].$$

Then there is an irreducible $\pi \in \widehat{G}$ *with* $m_\pi(M) \neq 0$ *and the following additional property. Pick a representative* $\xi \in (\mathfrak{h}^f_{\mathbb{C}})^*$ *for the infinitesimal character of* π. *Then the canonical real part of* ξ *(Definition 8.5) is at least as large as* $\| [N] \|$:

$$\| \operatorname{RE} \xi \| \geq \| [N] \|.$$

The kind of property that we want to prove is this: if the support of N is small, then $\| [N] \|$ must be large. Here is the formal setting.

Definition 8.8. Suppose $\mathcal{B}$ is a closed $K_{\mathbb{C}}$-invariant subset of $\mathcal{N}^*_\theta$. (The notation is chosen because we will most often want $\mathcal{B}$ to be the boundary $\overline{S} - S$ of a single orbit S.) Consider the subcategory of $\mathcal{C}(\mathfrak{g}, K)$ consisting of modules supported on $\mathcal{B}$:

$$\mathcal{C}(\mathfrak{g}, K)(\mathcal{B}) = \{M \in \mathcal{C}(\mathfrak{g}, K) \mid \mathcal{V}(M) \subset \mathcal{B}\}.$$

The Grothendieck group of this category is the subgroup of $K\mathcal{C}(\mathfrak{g}, K)$ spanned by the various $[N(\lambda_i, \tau)]$ (Proposition 7.9) with $\lambda_i \in \mathcal{B}$.

Set

$$\|\mathcal{B}\| = \min_{0 \neq [N] \in K\mathcal{C}(\mathfrak{g}, K)(\mathcal{B})} \| [N] \|,$$

the *infinitesimal character size of* $\mathcal{B}$.

Corollary 8.9. *Suppose* M *is an irreducible Harish-Chandra module for* G, *and* $\xi \in (\mathfrak{h}^f_{\mathbb{C}})^*$ *represents the infinitesimal character of* M. *Then the canonical real part of* ξ *is at least as large as the infinitesimal character size of* $\mathcal{V}(\operatorname{gr} M)$.

Notice that the infinitesimal character size of $\mathcal{B}$ is defined in terms of the change of basis matrix from Proposition 7.9 to Theorem 8.2: it is really a commutative algebra object. The Corollary gets interesting representation-theoretic information about M (a lower bound on its infinitesimal character) from this commutative algebra calculation.

In order to say something about unitarity, we need some ideas from [**VoCI**].

Definition 8.10. Suppose $M \in \mathcal{M}(\mathfrak{g}, K)$ has an invariant Hermitian form $\langle , \rangle_M$. Write $m_\mu(M)$ for the multiplicity in M of $\mu \in \widehat{K}$ (Definition 6.1). As is explained in [**VoCI**], the signature of the form gives rise to three non-negative integers

$$p_\mu(M), \quad q_\mu(M), \quad z_\mu(M), \qquad p_\mu(M) + q_\mu(M) + z_\mu(M) = m_\mu(M).$$

The form is non-degenerate if and only if $z(M) = 0$, positive semidefinite if and only if $q(M) = 0$, and so on. We can therefore define genuine virtual representations

$$[M]^+, \quad [M]^-, \quad [M]^0 \in K\mathcal{A}(K),$$

by

$$[M]^+ = \sum_{\mu \in \widehat{K}} p_\mu(M)[V_\mu], \qquad [M]^- = \sum_{\mu \in \widehat{K}} q_\mu(M)[V_\mu],$$

and so on.

Theorem 8.11 ([**VoCI**]). *Suppose $M \in \mathcal{M}(\mathfrak{g}, K)$ has an invariant Hermitian form; use the notation of Definition 8.10. The virtual K representations $[M]^\pm$ and $[M]^0$ are (restrictions to K of) virtual tempered representations of G of real infinitesimal character. According to Theorem 8.2, we may therefore regard them as well-defined virtual equivariant modules on the nilpotent cone:*

$$[M]^+, \quad [M]^-, \quad [M]^0 \in K\mathcal{C}(\mathfrak{g}, K).$$

Each infinitesimal character for a tempered representation appearing in $[M]^\pm$ or $[M]^0$ is bounded above by the canonical real part of the infinitesimal character of an irreducible constituent of M. In particular, if M has infinitesimal character ξ, then

$$\| \mathrm{RE}\, \xi \| \geq \| [M]^\pm \|.$$

Notice that the virtual modules $[M]^\pm$ and $[M]^0$ all have non-negative K-multiplicities, and their sum is the genuine virtual module $[\mathrm{gr}\, M]$. It is therefore natural to expect that $[M]^\pm$ and $[M]^0$ are themselves genuine virtual modules. It is very likely that this follows from the proof of Theorem 8.11, but I have not yet checked carefully. I will therefore be conservative and label it as a conjecture.

Conjecture 8.12. In the setting of Theorem 8.11, the virtual modules $[M]^\pm$ and $[M]^0$ are represented by actual modules $N^\pm$ and N^0 in $\mathcal{C}(\mathfrak{g}, K)$.

At last we can use these ideas to say something about unitary representations.

Theorem 8.13. *Suppose M is a $(\mathfrak{g}, K)$ module of finite length and infinitesimal character ξ, carrying a non-degenerate invariant Hermitian form $\langle , \rangle_M$. Assume that*

1. *The associated variety $\mathcal{V}(\mathrm{gr}\, M)$ is the closure of single orbit $S = K_{\mathbb{C}} \cdot \lambda$. Write*
$$\mathcal{B} = \overline{S} - S$$
for the boundary of this orbit.
2. *The genuine virtual representation $\tau = \chi(\lambda, \mathrm{gr}\, M)$ defined by Proposition 7.6 is irreducible.*
3. *The canonical real part of the infinitesimal character of M is strictly smaller than the infinitesimal character size of $\mathcal{B}$ (Definition 8.8):*
$$\| \mathrm{RE}\, \xi \| < \| \mathcal{B} \|.$$
4. *Conjecture 8.12 holds for M.*

Then M is irreducible and the Hermitian form on M is definite, so M is the Harish-Chandra module of an irreducible unitary representation of G.

Suppose in addition that the module $N(\lambda, \tau)$ of Proposition 7.9 has infinitesimal character size less than $\|\mathcal{B}\|$:

$$\| [N(\lambda, \tau)] \| < \|\mathcal{B}\|.$$

Then $[\operatorname{gr} M] = [N(\lambda, \tau)]$. In particular, if τ is an admissible orbit datum at λ (Definition 5.2), then M satisfies the desideratum at (5.5) to be a quantization of the corresponding real nilpotent coadjoint orbit.

Proof. Fix modules $N^{\pm}$ as in Conjecture 8.12. (Since the form is non-degenerate, the radical is zero, so $N^0 = 0$.) Because these sum to $\operatorname{gr} M$ in the Grothendieck group, it follows that $\mathcal{V}(N^{\pm}) \subset \mathcal{V}(\operatorname{gr} M)$. We therefore have well-defined genuine virtual representations $\chi(\lambda, N^{\pm})$, and

$$\chi(\lambda, \operatorname{gr} M) = \chi(\lambda, N^+) + \chi(\lambda, N^-).$$

On the other hand, assumption (2) in the theorem says that the left side is irreducible. It follows that one of the terms on the right is zero. Possibly after replacing $\langle , \rangle_M$ by its negative (which interchanges N^+ and N^-), we conclude that

$$\chi(\lambda, N^+) = \chi(\lambda, \operatorname{gr} M), \qquad \chi(\lambda, N^-) = 0.$$

By Proposition 7.6, it follows that $\mathcal{V}(N^-) \subset \mathcal{B}$.

We are trying to prove that $\langle , \rangle_M$ is positive definite. Suppose not; that is, that $N^- \neq 0$. By Definition 8.8 and what we have just seen about the support of N^-,

$$\| [N^-] \| \geq \|\mathcal{B}\|.$$

According to Theorem 8.11,

$$\| \operatorname{RE} \xi \| \geq \| [N^-] \|.$$

These two inequalites together contradict the assumption in (3) of the theorem. The contradiction proves the positivity of the Hermitian form. A similar argument proves that M must be irreducible.

Finally, we consider the claim that $[\operatorname{gr} M] = [N(\lambda, \tau)]$. According to Propositions 7.6 and 7.9, the difference $[\operatorname{gr} M] - [N(\lambda, \tau)]$ must be in the Grothendieck group of modules supported on $\mathcal{B}$. Suppose it is non-zero. Assumption (3) in the theorem and the assumption about the infinitesimal character size of $N(\lambda, \tau)$ imply that

$$\| [\operatorname{gr} M] - [N(\lambda, \tau)] \| < \|\mathcal{B}\|,$$

contradicting Definition 8.8. The virtual module must therefore be zero, as we wished to show. ☐

Theorem 8.13 produces irreducible unitary representations attached to a nilpotent coadjoint orbit; so we should consider how to go about verifying its hypotheses.

Arranging condition (1) is relatively straightforward; here is a very brief sketch. Begin with the corresponding complex group orbit

$$S_{\mathbb{C}} = \mathrm{Ad}(\mathfrak{g}_{\mathbb{C}}) \cdot \lambda \subset \mathfrak{g}_{\mathbb{C}}^*; \tag{8.14(a)}$$

here $\mathrm{Ad}(\mathfrak{g}_{\mathbb{C}})$ is the identity component of the automorphism group of $\mathfrak{g}_{\mathbb{C}}$. Choose an infinitesimal character ξ (cleverly!) in such a way that the maximal primitive ideal J_ξ^{max} of infinitesimal character ξ has associated variety $\overline{S_{\mathbb{C}}}$. (Methods for choosing such a ξ may be found in [**BV**].) Finally let M be any irreducible Harish-Chandra module with annihilator J_ξ^{max}. (Methods for constructing such modules in the case of complex groups may also be found in [**BV**]; similar ideas apply to real groups.) Then $\mathcal{V}(\operatorname{gr} M)$ will be the closure of the union of some of the nilpotent $K_{\mathbb{C}}$ orbits in $S_{\mathbb{C}} \cap (\mathfrak{g}/\mathfrak{k})^*$. This is close to (1).

The irreducibility condition in (2) of Theorem 8.13 is a subtle point, and I will not try to say more about it here; but the methods described below for studying (3) are helpful.

For condition (3), we already know ξ explicitly from arranging (1); so the whole problem is to find an explicit lower bound $\|\mathcal{B}\| \geq C$ (Definition 8.8). This problem has two parts. The modules involved in the definition of $\|\mathcal{B}\|$ are integer combinations of various $N(\lambda', \tau')$ (Proposition 7.9) with $\lambda' \in \mathcal{B}$. The first part of the problem is to find more or less explicit formulas

$$[N(\lambda', \tau')] = \sum_{\mu \in \widehat{K}} m_\mu(\lambda', \tau')[M(\mu)]. \tag{8.14(b)}$$

This is an algebraic geometry problem; it can be approached by the technique McGovern introduced for studying functions on nilpotent orbits. One introduction to this technique is in [**AHV**]. In particular, one hopes to find in each of these formulas a term $M(\mu)$ with $m_\mu(\lambda', \tau') \neq 0$ and

$$\|\lambda(\mu)\| \geq C.$$

The second part of the problem is to check that some of these large terms $M(\mu)$ do not cancel when we form integer combinations

$$\sum_{\lambda' \in \mathcal{B},\ \tau' \in \widehat{K^{\lambda'}}} n_{\lambda', \tau'}[N(\lambda', \tau')]. \tag{8.14(c)}$$

Once we know the formulas in (8.14)(b), this is a question of linear algebra.

Finally, I hope that verification of (4) is just a matter of examining [**VoCI**].

To conclude, here is the simplest example of a computation of $\|\mathcal{B}\|$. This result controls the infinitesimal character of a virtual Harish-Chandra module with finite K-multiplicities. The most obvious way to produce such a virtual module is as an integer combination of finite-dimensional Harish-Chandra modules. There is an obvious lower bound for the infinitesimal character of such a module. The theorem says that more exotic methods of producing the modules cannot give smaller infinitesimal character.

Theorem 8.15. *Suppose $\mathcal{B} = \{0\} \subset (\mathfrak{g}/\mathfrak{k})^*$. Then the infinitesimal character size of $\mathcal{B}$ (Definition 8.7) is equal to the length of ρ, half the sum of a set of positive roots.*

The trivial representation M of G has infinitesimal character ρ and associated variety $\{0\}$, so the inequality $\mathcal{B} \leq \|\rho\|$ follows from Corollary 8.9. We consider therefore the opposite inequality.

Proof for G complex connected. The objects $[N]$ appearing in Definition 8.8 are integer combinations of $N(0,\tau)$, with τ an irreducible representation of the isotropy group $K_{\mathbb{C}}$ of 0. By inspection of Proposition 7.9, $N(0,\tau)$ is just the irreducible representation (τ, V_τ) of K, equipped with the module structure in which a polynomial $p \in S(\mathfrak{g}/\mathfrak{k})$ acts by its value at 0. We need to express $N(0,\tau)$ (as a representation of K) as an integer combination of the tempered representations $M(\mu)$.

In Theorem 8.1 those representations are parametrized by irreducible representations μ of K; but now it will be more convenient to parametrize them by ($W(K,T)$ group orbits of) elements of $\widehat{T}$ (Proposition 8.4). Write $H^f = TA$ for the corresponding Cartan in G, and $B = TAN$ for a Borel subgroup. Then the tempered representation corresponding to $\gamma \in \widehat{T}$ is

$$M(\gamma) = \mathrm{Ind}_B^G(\gamma \otimes 1 \otimes 1); \qquad M(\gamma)|_K = \mathrm{Ind}_T^K(\gamma). \tag{8.16(a)}$$

Write

$$\Delta^+(\mathfrak{k}_{\mathbb{C}}, \mathfrak{t}_{\mathbb{C}}) \subset \widehat{T} \tag{8.16(b)}$$

for a set of positive roots of T in $\mathfrak{k}_{\mathbb{C}}$, and

$$2\rho_c = \sum_{\alpha \in \Delta^+(\mathfrak{k}_{\mathbb{C}}, \mathfrak{t}_{\mathbb{C}})} \alpha. \tag{8.16(c)}$$

Then the differential of $2\rho_c$ is equal to ρ (for a certain choice of positive roots of $\mathfrak{h}^f_{\mathbb{C}}$); so the inequality we are trying to prove may be written as

$$\|\mathcal{B}\| \geq \|2\rho_c\|. \tag{8.16(d)}$$

Write $\gamma_\tau \in \widehat{T}$ for the highest weight of τ. A version of the Weyl character formula is

$$\tau = \sum_{S \subset \Delta^+(\mathfrak{k}_{\mathbb{C}}, \mathfrak{t}_{\mathbb{C}})} (-1)^{|S|}\, \mathrm{Ind}_T^K(\gamma_\tau + 2\rho(S)). \tag{8.16(e)}$$

Here $2\rho(S)$ is the sum of the roots in S. Combining this with (8.16)(a) gives

$$[N(0,\tau)] = \sum_{S \subset \Delta^+(\mathfrak{k}_{\mathbb{C}}, \mathfrak{t}_{\mathbb{C}})} (-1)^{|S|} M(\gamma_\tau + 2\rho(S)). \tag{8.16(f)}$$

Recall from the remark after Proposition 8.4 that $\lambda(\gamma)$ is just the differential of γ. It is a standard and elementary fact about compact group representations that the longest of the weights appearing in (8.16)(f) is $\gamma_\tau + 2\rho_c$:

$$[N(0,\tau)] = \pm M(\gamma_\tau + 2\rho_c) + \text{terms } M(\gamma') \text{ with } \lambda(\gamma') \text{ strictly shorter.} \tag{8.16(g)}$$

Suppose now that $N = \sum_\tau m_\tau [N(0,\tau)]$ is a non-zero finite integer combination as in Definition 8.8. Choose τ_0 so that $\|\gamma_{\tau_0} + 2\rho_c\|$ is maximal subject to $m_\tau \neq 0$. It follows from (8.16)(g) that

$$\|\,[N]\,\| = \|\gamma_{\tau_0} + 2\rho_c\| \geq \|2\rho_c\|. \tag{8.16(h)}$$

This proves (8.16)(d), which is what we wished to show. □

EXERCISES

Exercise 1. Find all the Riemannian homogeneous spaces for $SO(3)$. (Decide first of all what it should mean for two such spaces to be isomorphic.)

Exercise 2. Show that every symplectic homogeneous space for $SO(3)$ is either a point, or a two-dimensional sphere of radius $r > 0$. Consequently these spaces are parametrized by the non-negative real numbers.

Exercise 3. This exercise concerns the non-existence of a nice simultaneous quantization of a large family of classical observables. The results come more or less from [**Gr**], with a more mathematical account in [**vH**]. Perhaps the best place to read about it is in [**AM**], Theorem 5.4.9. We take as symplectic manifold $X = \mathbb{R}^2$, with coordinate functions p and q. The Poisson bracket on smooth functions on X is

$$\{f, g\} = \frac{\partial f}{\partial p}\frac{\partial g}{\partial q} - \frac{\partial f}{\partial q}\frac{\partial g}{\partial p}. \qquad E3(a)$$

For the quantization, we use the Hilbert space $\mathcal{H} = L^2(\mathbb{R})$; the coordinate on $\mathbb{R}$ is called q. We begin the quantization as in (1.6), defining two skew-adjoint operators

$$A(q) = \text{multiplication by } \sqrt{-1}q, \qquad A(p) = \partial/\partial q. \qquad E3(b)$$

Recall that the quantum interpretation of position suggests defining

$$A(f) = \text{multiplication by } \sqrt{-1}f, \qquad f \in \mathbb{R}[q]. \qquad E3(c)$$

As in (1.7), we try to quantize other polynomial functions on $\mathbb{R}^2$ using polynomial coefficient differential operators on $\mathbb{R}$. The requirement (1.4) is

$$[A(f), A(g)] = A(\{f, g\}) \qquad (f, g \in \mathbb{R}[p, q]). \qquad E3(d)$$

It is easy to check that this is satisfied for those functions where A has so far been defined: polynomials in q, and linear functions of p. On quadratic polynomials, it is natural to define

$$A(p^2) = \sqrt{-1}\frac{\partial^2}{\partial q^2}, \quad A(pq) = -\frac{1}{2}(q\frac{\partial}{\partial q} + \frac{\partial}{\partial q}q), \quad A(q^2) = \sqrt{-1}q^2. \qquad E3(e)$$

(We have already explained $A(q^2)$. The requirement (E3)(d) means that $A(p^2)$ should commute with $A(p) = \partial/\partial q$, so it should be a constant coefficient differential operator. This particular choice seems reasonable, and is consistent with the standard construction of the quantum harmonic oscillator. Finally the choice of $A(pq)$ is forced by (E3)(d), since $\{p^2, q^2\} = 4pq$.)

Here at last is the exercise. Show first that the span of p^2, q^2, and pq is a Lie algebra $\mathfrak{g}$ under Poisson bracket, isomorphic to $SL(2,\mathbb{R})$. Next, show that the Poisson bracket defines a representation of $\mathfrak{g}$ on the space S_n of homogeneous polynomials of degree n in p and q; and that this is the irreducible $n+1$-dimensional representation of $\mathfrak{g}$.

Now write T_n for the space of skew-symmetric polynomial coefficient differential operators on $\mathbb{R}$, of total degree at most n. (The total degree is the sum of the differential operator degree and the polynomial coefficient degree.) Show that T_n is invariant under operator bracket by $A(\mathfrak{g})$ (the operators of (E3)(e)), and that this defines a representation of $\mathfrak{g}$ on T_n. Show further that T_n is the sum of the irreducible representations of $\mathfrak{g}$ of dimensions 1, 2, ... , $n+1$. Show that $A(q^n)$ may be regarded as a highest weight vector for the representation of dimension $n+1$.

Conclude that there is a unique linear "quantization map" A carrying $\mathbb{R}[p,q]$ to skew-symmetric polynomial coefficient differential operators on $\mathbb{R}$, subject to the requirements (E3)(e); $A(q^n) = \sqrt{-1}q^n$; and

$$A(\{f,g\}) = [A(f), A(g)] \qquad (f \text{ homogeneous of degree } 2). \qquad E3(f)$$

What remains is to show that A fails to satisfy (E3)(d) in general. To do that, show that $\{p^2, \{p^2, q^4\}\}$ is proportional to $\{p^3, q^3\}$, but that $[A(p^2), [A(p^2), A(q^4)]]$ is not proportional to $[A(p^3), A(q^3)]$.

Exercise 4. Suppose ϕ is a non-degenerate bilinear form on $\mathbb{R}^n$, either symmetric or skew-symmetric. There is a unique invertible matrix J so that

$$\phi(u,v) = \langle Ju, v\rangle \qquad (u, v \in \mathbb{R}^n).$$

Assume that $J^2 = \pm I$. Show that the group of the form

$$G(\phi) = \{g \in GL(n,\mathbb{R}) \mid \phi(gu, gv) = \phi(u,v) \quad (u, v \in \mathbb{R}^n)\}$$

is a linear reductive group. Check that this condition allows the indefinite orthogonal groups $O(p,q)$ $(p+q=n)$ and the symplectic group $Sp(2n,\mathbb{R})$ as linear reductive groups.

Exercise 5. Find a closed subgroup of $GL(1,\mathbb{R})$ that is stable under θ but is not a reductive group. (Hint: it follows from Theorem 2.6 that a reductive group has a finite number of connected components.) Can you find more interesting examples in $GL(2,\mathbb{R})$?

Exercise 6. Suppose G is a linear reductive Lie group (Definition 2.5), and we are given elements H, E, and F of the Lie algebra as in (2.9)(a). Why does the group homomorphism ϕ of (2.9)(b–c) exist? (Hint: the group $SL(2,\mathbb{R})$ is not simply connected, so the problem is not completely trivial.) What can you say if G is only assumed to be reductive?

Exercise 7. Write $GL(n, \mathbb{C}) = K \cdot \exp(\mathfrak{s})$ for the polar decomposition of $n \times n$ complex matrices; here $K = U(n)$ and $\mathfrak{s}$ is the space of self-adjoint matrices. Fix $X \in \mathfrak{k}$. Show that the set of fixed points of the one-parameter group $\mathrm{Ad}(\exp(tX))$ on $\mathfrak{s}$ is precisely $\mathfrak{s}^X$. Conclude that the stabilizer of X in the adjoint action of $GL(n, \mathbb{C})$ is $K^X \cdot \exp(\mathfrak{s}^X)$.

Suppose that $Z \in \mathfrak{s}$. Show that the stabilizer of Z in the adjoint action is $K^Z \cdot \exp(\mathfrak{s}^Z)$. (Hint: look at iZ.)

Suppose $G = K \cdot \exp(\mathfrak{s})$ is the Cartan decomposition of a reductive group, and $X \in \mathfrak{s}$. Show that $G^X = K^X \cdot \exp(\mathfrak{s}^X)$, and conclude that G^X is a reductive group.

Exercise 8. Suppose G is a reductive group with Cartan involution θ.

1. Show that the bilinear form Q on $\mathfrak{g}$ defined by

$$Q(X, Y) = -\langle X, \theta Y \rangle$$

(notation as in Proposition 2.7) is positive definite.
2. If $X \in \mathfrak{s}$, show that the operator $\mathrm{ad}(X)$ on $\mathfrak{g}$ is self-adjoint with respect to Q. Conclude that $\mathrm{ad}(X)$ is diagonalizable with real eigenvalues.
3. If $Y \in \mathfrak{k}$, show that the operator $\mathrm{ad}(Y)$ on $\mathfrak{g}$ is skew-adjoint with respect to Q. Conclude that $\mathrm{ad}(Y)_{\mathbb{C}}$ is diagonalizable with purely imaginary eigenvalues.
4. If $k \in K$, show that the operator $\mathrm{Ad}(k)$ on $\mathfrak{g}$ is orthogonal with respect to Q. Conclude that $\mathrm{Ad}(k)_{\mathbb{C}}$ is diagonalizable with eigenvalues in the unit circle.

Exercise 9. The point of this exercise is to understand the constructions of (3.2) in some examples, and in particular the character $2\rho_{\mathfrak{q}}$ of Proposition 3.3. You can pick your own favorite examples; certainly even compact groups are interesting in this context. If you would like to test your grasp of general structure theory, you can try to prove that if G is a complex reductive group (and λ_e is elliptic, and so on) then the character $2\rho_{\mathfrak{q}}$ always has a distinguished square root.

If you don't have a favorite example to try, here is one. Begin with the complex vector space $\mathbb{C}^n$. Using a standard identification $\mathbb{C} \simeq \mathbb{R}^2$, we get $\mathbb{C}^n \simeq \mathbb{R}^{2n}$. Let $G = GL(2n, \mathbb{R})$, so that $\mathfrak{g}$ consists of all $\mathbb{R}$-linear transformations of $\mathbb{R}^{2n}$. Define

$$J = \text{multiplication by } i \text{ on } \mathbb{C}^n \simeq \mathbb{R}^{2n}.$$

This is an $\mathbb{R}$-linear transformation, and so an element of $\mathfrak{g}$; it is elliptic. (Why?) Write $\lambda_e \in \mathfrak{g}^*$ for the corresponding linear functional:

$$\lambda_e(Z) = \operatorname{tr} JZ.$$

Use the notation of (3.2). In particular L consists of invertible real-linear transformations of $\mathbb{C}^n$ that commute with multiplication by i; so

$$L = GL(n, \mathbb{C}).$$

Having come this far, you should try to identify the coadjoint orbit $G \cdot \lambda_e$ with the space of all complex structures on the vector space $\mathbb{R}^{2n}$.

Now we want to investigate the eigenspaces of $i\,\mathrm{ad}(J)$. Let B (for "bar") be the complex-conjugation linear transformation of $\mathbb{C}^n$. Then $JB = -BJ$. Show that $B + iJB$ is in the $+2$ eigenspace of $i\,\mathrm{ad}(J)$, and in fact that

$$\mathfrak{u} = \{T(B + iJB) \mid T \text{ a complex-linear transformation of } \mathbb{C}^n\}.$$

Deduce that the adjoint action of L on $\mathfrak{u}$ is equivalent to the action of $GL(n, \mathbb{C})$ on $n \times n$ complex matrices, by

$$g \cdot T = gT\overline{g}^{-1} \qquad (g \in GL(n, \mathbb{C}), T \in \mathfrak{gl}(n, \mathbb{C})).$$

Finally conclude that

$$2\rho_{\mathfrak{q}}(g) = (\det(g)/\overline{\det(g)})^n \qquad (g \in GL(n, \mathbb{C})).$$

Does this unitary character of $GL(n, \mathbb{C})$ have a square root?

Exercise 10. One of the pleasant properties of a Stein manifold X is that a compact complex submanifold of X must be finite. In the setting of (3.2), show that $K/L \cap K$ is a compact complex submanifold of X. Deduce that X can be Stein only if the Lie algebra element X_e belongs to the center of $\mathfrak{k}$. (In that case it turns out that X actually is Stein.)

If you know something about the Bott-Borel-Weil theorem, show that the holomorphic vector bundle $\mathcal{V}_\tau$ has non-zero holomorphic sections over $K/L \cap K$ if and only if $K/L \cap K$ is finite. Deduce that $\mathcal{V}_\tau$ has non-zero holomorphic sections over X only if $K/L \cap K$ is finite. (This statement is also if and only if.)

Exercise 11. Use the notation of Definition 5.2. Find a definition of a *metaplectic double cover of* $K_{\mathbb{C}}^{\lambda_\theta}$ analogous to (3.6), and a character ρ of $\widetilde{K_{\mathbb{C}}^{\lambda_\theta}}$. Show that admissible orbit data at λ_θ are in one-to-one correspondence with irreducible genuine representations $(\widetilde{\tau}, V_{\widetilde{\tau}})$ of this metaplectic cover that are trivial on the identity component (that is, having differential 0). To complete the analogy with Proposition 3.7, show that the restriction of λ_θ to $\mathfrak{k}_{\mathbb{C}}^{\lambda_\theta}$ is zero.

Exercise 12. Suppose $G = Sp(2n, \mathbb{R})$ is the standard real symplectic group. We want to look at some examples of nilpotent admissible orbit data. The facts stated in this first paragraph can just be assumed, although verifying them makes a reasonable exercise in structure theory. The maximal compact subgroup of G is isomorphic to $U(n)$, so its complexification $K_{\mathbb{C}}$ is isomorphic to $GL(n, \mathbb{C})$. There are two representations of $K_{\mathbb{C}}$ on the space S of symmetric $n \times n$ complex matrices: we write them as (π^+, S^+) and (π^-, S^-). Explicitly,

$$\pi^+(g)U = gU\,{}^tg \qquad (g \in GL(n, \mathbb{C}), U \in S)$$

and

$$\pi^-(g)V = {}^tg^{-1}Vg^{-1} \qquad (g \in GL(n, \mathbb{C}), V \in S)$$

Then there is a $K_{\mathbb{C}}$-equivariant isomorphism

$$\mathfrak{s}_{\mathbb{C}}^* \simeq S^+ \oplus S^-.$$

The nilpotent elements of $\mathfrak{s}_{\mathbb{C}}^*$ are the pairs (U,V) so that UV and VU are both nilpotent. (A complete classification of the $K_{\mathbb{C}}$ orbits of such pairs is an entertaining linear algebra problem, which the reader might also like to consider.) We will confine our attention to certain relatively small orbits: specifically, to the case $UV = VU = 0$. There are two obvious invariants to attach to such an orbit: the rank p of U and the rank q of V. The vanishing of UV forces $p+q \le n$. It turns out that the pair (p,q) determines the $K_{\mathbb{C}}$ orbit of (U,V).

So let us fix (p,q) with $p+q \le n$. Define

$$U = \begin{pmatrix} I_p & 0 \\ 0 & 0_{n-p} \end{pmatrix}, \qquad V = \begin{pmatrix} 0_{n-q} & 0 \\ 0 & I_q \end{pmatrix}, \qquad \lambda_\theta = (U,V).$$

Then λ_θ is a nilpotent element of $\mathfrak{s}_{\mathbb{C}}^*$ in the orbit parametrized by (p,q). We want to compute the admissible orbit data. Show first of all that the stabilizer of λ_θ in $K_{\mathbb{C}}$ is

$$K_{\mathbb{C}}^{\lambda_\theta} = \begin{pmatrix} A & * & * \\ 0 & B & * \\ 0 & 0 & C \end{pmatrix} \qquad (A \in O(p,\mathbb{C}), B \in GL(n-p-q,\mathbb{C}), C \in O(q,\mathbb{C})).$$

Next, show that the character 2ρ of Definition 5.2 is

$$2\rho(A,B,C) = (\det A)^{n-1}(\det B)^{q-p}(\det C)^{n-1}.$$

(Here we write (A,B,C) as shorthand for the matrix from the preceding display.) Conclude that admissible orbit data exist if and only if either $q-p$ is even, or $p+q=n$. If one of these conditions is satisfied, then the possible admissible orbit data are

$$\tau(A,B,C) = (\det A)^{\epsilon}(\det B)^{(q-p)/2}(\det B)^{\delta}.$$

Here ϵ and δ both belong to $\mathbb{Z}/2\mathbb{Z}$; the factor with ϵ (respectively δ) actually appears only if p (respectively q) is non-zero. Therefore (still assuming that $q-p$ is even, or that $p+q=n$) there are four admissible orbit data if p and q are both non-zero; two if exactly one of p and q is zero; and one if $p=q=0$.

Exercise 13. This is a continuation of Exercise 12. First show that the orbit with parameters (p,q) has complex dimension $(p+q)(p+q+1)/2+(p+q)(n-p-q)$. The orbits in the boundary are those with parameters (p',q'), with $p' \le p$, $q' \le q$, and $p'+q' < p+q$. Conclude that the boundary has complex codimension $n-(p+q)+1$ (if $p+q>0$; the boundary is empty if $p+q=0$). Therefore Definition 5.4 applies if and only if $p+q<n$. If you know a little about the representations of $U(n)$, you might try to calculate $X_K(\lambda_n,\tau)$ in this case. It turns out that the representations of $U(n)$ appear with multiplicity one; those appearing are the ones of highest weight

$$\mu = (\mu_1, \dots, \mu_q, (q-p)/2, \dots, (q-p)/2, \mu_{n-p+1}, \dots, \mu_n),$$

$$\mu_i \equiv \delta \pmod 2 \quad (1 \le i \le q), \qquad \mu_j \equiv \epsilon \pmod 2 \quad (n-p+1 \le j \le n).$$

Of course μ must also be a highest weight: that is, the coordinates of μ must be weakly decreasing integers. (Hint: the most difficult part is to understand induction from $O(p)$ to $U(p)$. This is a compact symmetric space, so the induction is computed by a theorem of Helgason. It is the sum of all representations of $U(p)$ whose highest weights have all coordinates even. Given this fact and its twist by the determinant character, you just need the Borel-Weil theorem to finish the calculation.)

Exercise 14. This exercise looks at the construction of Proposition 7.9 in the case of $G = U(1,1)$. Then $K = U(1) \times U(1)$, and $K_{\mathbb{C}} = \mathbb{C}^\times \times \mathbb{C}^\times$. We have

$$(\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}})^* \simeq \left\{ \begin{pmatrix} 0 & u \\ v & 0 \end{pmatrix} \mid u, v \in \mathbb{C} \right\} \simeq \mathbb{C}^2;$$

the action of $K_{\mathbb{C}}$ is

$$(z, w) \cdot (u, v) = (zw^{-1}u, z^{-1}wv).$$

The nilpotent cone is

$$\mathcal{N}_\theta^* = \{(u, v) \mid uv = 0\}.$$

There are three orbits of $K_{\mathbb{C}}$, represented by

$$\lambda_1 = (1, 0), \qquad \lambda_2 = (0, 1), \qquad \lambda_3 = (0, 0).$$

The first two isotropy subgroups are

$$H_1 = H_2 = \mathbb{C}_\Delta^\times \subset \mathbb{C}^\times \times \mathbb{C}^\times.$$

The third is $H_3 = K_{\mathbb{C}}$.

Fix an integer m, and let τ_m be the character

$$\tau_m(z) = z^m$$

of H_1. Then $M(\lambda_1, \tau_m)$ may be identified with the space of functions on $K_{\mathbb{C}}$ transforming according to the character τ_m of H_1. As a basis for $M(\lambda_1, \tau_m)$, we can take the functions

$$f_p(z, w) = z^p w^{m-p}.$$

The function f_p transforms under K according to the character $(p, m-p)$ (in the standard identification of $\widehat{K}$ with $\mathbb{Z}^2$).

Now $M(\lambda_1, \tau_m)$ is also a module for $S(\mathfrak{g}/\mathfrak{k}) \simeq \mathbb{C}[u, v]$. Show that in this module structure, v acts by zero, and $u \cdot f_p = f_{p+1}$. Conclude that $M(\lambda_1, \tau_m)$ is not finitely generated. Show that the module $N(\lambda_1, \tau_m)$ constructed in the proof of Proposition 7.9 must be

$$N(\lambda_1, \tau_m) = \text{span of } \{f_p \mid p \geq p_0\}$$

for some integer p_0. Conclude that

$$K\text{-types of } \mathrm{Ind}_{H_1}^{K_{\mathbb{C}}}(\tau_m) = \{(r, s) \mid r + s = m\}$$

$$K\text{-types of } N(\lambda_1, \tau_m) = \{(r, s) \mid r + s = m,\ r - s \geq 2p_0 - m\}.$$

If you are feeling very ambitious, much of this exericise can be generalized to the setting of Exercise 12, with λ_1 one of the elements with $p + q = n$.

Exercise 15. This exercise looks at the bases provided by Theorem 8.2 and Proposition 7.9 in case $G = SL(2, \mathbb{C})$ (viewed as a real group). The maximal compact subgroup is $K = SU(2)$. Let T be the standard maximal torus in K (consisting of diagonal matrices); T is isomorphic to $U(1)$, so $\widehat{T} \simeq \mathbb{Z}$; we write δ_m for the character corresponding to m. The representations of K are parametrized (by their highest weights) by non-negative integers; we write μ_m for the ($m+1$-dimensional) representation of K of highest weight δ_m.

Let A be the group of diagonal matrices in G with positive real entries, and N the group of upper triangular matrices with ones on the diagonal. Then $B = TAN$ is a Borel subgroup of G. If $\delta \in \widehat{T}$ and $\nu \in \widehat{A}$, then we can construct a principal series representation

$$I(\delta \otimes \nu) = \mathrm{Ind}_B^G(\delta \otimes \nu \otimes 1).$$

This representation is tempered exactly when ν is unitary, and has real infinitesimal character exactly when ν is real-valued. We will often think of $\widehat{A}$ as identified with $\mathfrak{a}^*_{\mathbb{C}}$, and so write $-\nu$ for the inverse of the character ν. The principal series $I(\delta \otimes \nu)$ and $I(\delta^{-1} \otimes -\nu)$ have the same irreducible composition factors, so for many purposes we can confine our attention to characters $\delta = \delta_m$ of T with $m \geq 0$.

Because of the Iwasawa decomposition $G = KAN$, we find

$$I(\delta \otimes \nu)|_K \simeq \mathrm{Ind}_T^K(\delta),$$

and this in turn is the sum of all the irreducible representations of K containing the weight δ. For $m \geq 0$, this is

$$I(\delta_m \otimes \nu)|_K \simeq \sum_{k \geq 0} \mu_{m+2k}.$$

The tempered irreducible representations of real infinitesimal character are the various principal series

$$\pi_m = I(\delta_m \otimes 0) \qquad (m \geq 0).$$

Their restrictions to K are what we just computed:

$$\pi_m|_K \simeq \sum_{k \geq 0} \mu_{m+2k}.$$

These representations are the basis of Theorem 8.2 for (restrictions to K of) virtual Harish-Chandra modules. That is, any finite length Harish-Chandra module for G has the same restriction to K as a unique integer combination of the various π_m. The first part of the exercise is essentially to verify this assertion. There are just two kinds of irreducible Harish-Chandra modules for G: the irreducible principal series, and the finite-dimensional representations. For the principal series, we have

$$I(\delta_m \otimes \nu)|_K \simeq \pi_{|m|}|_K.$$

Show that if F is any irreducible finite-dimensional representation of G, then

$$F|_K \simeq \mu_p + \mu_{p+2} + \cdots + \mu_{p+2q},$$

for some non-negative integers p and q. Conclude that

$$F|_K \simeq (\pi_p - \pi_{p+2q+2})|_K.$$

Now we examine the basis of Proposition 7.9. The complexification of K is G, and the action of $K_{\mathbb{C}}$ on $(\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}})^*$ is equivalent to the action of G on $\mathfrak{g}$. In particular there are exactly two nilpotent orbits, through $\lambda_1 = 0$ and $\lambda_2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. The isotropy groups are $H_1 = G$ and $H_2 = Z(G)N$. The orbit dimensions are 0 and 2, so the codimension condition in Proposition 7.9 is always satisfied. The irreducible (algebraic) representations of G may be identified with the irreducible representations μ_m of K. The irreducible representations of H_2 are the irreducible characters of $Z(G) \simeq \mathbb{Z}/2\mathbb{Z}$; there is the trivial character τ_e and the non-trivial character τ_o. The restrictions to K of the various standard modules are computed by Proposition 7.9; check that they are

$$\begin{aligned} N(\lambda_1, \mu_m)|_K &\simeq \mu_m, \\ N(\lambda_2, \tau_e)|_K &\simeq \sum_{k\geq 0} \mu_{2k}, \\ N(\lambda_2, \tau_o)|_K &\simeq \sum_{k\geq 0} \mu_{2k+1}. \end{aligned}$$

Proposition 7.9 says that any Harish-Chandra module for G has the same restriction to K as a unique integer combination of the various $N(\lambda, \tau)$. The next part of the exercise is to verify that assertion. By the first part, it is enough to consider the various π_m. Show that if m is even and non-negative,

$$\pi_m|_K \simeq (N(\lambda_2, \tau_e) - N(\lambda_1, \mu_0) - N(\lambda_1, \mu_2) - \cdots - N(\lambda_1, \mu_{m-2}))|_K;$$

and similarly for m odd.

The point of section 8 is that the other change of basis is more interesting. Show that

$$N(\lambda_1, \mu_m)|_K = (\pi_m - \pi_{m+2})|_K,$$
$$N(\lambda_2, \tau_e)|_K = \pi_0|_K, \qquad N(\lambda_2, \tau_o)|_K = \pi_1|_K.$$

Exercise 16. This exercise concerns the notion of "infinitesimal character size" in Definition 8.8, for the group $SL(2, \mathbb{C})$. In Exercise 15 we identified the characters of the compact torus T with $\mathbb{Z}$. This identification extends to $i\mathfrak{t}^* \simeq \mathbb{R}$. Let us also normalize the invariant bilinear form on $\mathfrak{g}$ so that (after restriction, complexification, and dualization to $i\mathfrak{t}^*$) it agrees with the standard inner product on $\mathbb{R}$. Thus the representation μ_m of K has highest weight δ_m, which (according to the remark after Proposition 8.4) is mapped by the algorithm of [**VG**] to the weight $m \in \mathbb{R} \simeq i\mathfrak{t}^*$, which has norm m. Using these facts, and the calculation in Exercise 15, prove that

$$\|K_{\mathbb{C}} \cdot \lambda_1\| = \|2\| = 2, \qquad \|K_{\mathbb{C}} \cdot \lambda_2\| = \|0\| = 0.$$

Next, show that $2 \in i\mathfrak{t}^*$ is equal to ρ, half the sum of a certain system of positive roots for the fundamental Cartan. This is the infinitesimal character of the trivial representation of G. The conclusion (by Proposition 8.7) is that any virtual Harish-Chandra module $[M]$ with $[\mathrm{gr}\, M]$ finite-dimensional and non-zero must have an irreducible constituent of infinitesimal character at least as big as ρ.

BIBLIOGRAPHY

[AM] R. Abraham and J. Marsden, *Foundations of Mechanics*, Benjamin/Cummings, Reading, Massachusetts, 1978.

[AHV] J. Adams, J.-S. Huang, and D. Vogan, *Functions on the model orbit in E_8*, Represent. Theory **2** (1998), 224–263.

[Arn] V. Arnold, *Mathematical Methods of Classical Mechanics*, Springer-Verlag, New York-Heidelberg-Berlin, 1978.

[BV] D. Barbasch and D. Vogan, *Unipotent representations of complex semisimple Lie groups*, Ann. of Math. **121** (1985), 41–110.

[CM] D. Collingwood and W. McGovern, *Nilpotent Orbits in Semisimple Lie Algebras*, Van Nostrand Reinhold, New York, 1993.

[D] M. Duflo, *Théorie de Mackey pour les groupes de Lie algébriques*, Acta Math. **149** (1982), 153–213.

[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.

[Gr] H. J. Groenewold, *On the principles of elementary quantum mechanics*, Physica **12** (1946), 405–460.

[EGA] A. Grothendieck (rédigés avec la collaboration de J. Dieudonné), *Eléments de Géométrie Algébrique IV. Etude locale des schémas et des morphismes de schémas (Seconde Partie)*, Publications Mathématiques, vol. 24, Institut des Hautes Etudes Scientifiques, Le Bois-Marie, Bures-sur-Yvette, France, 1965.

[vH] L. van Hove, *Sur certaines représentations unitaires d'un groupe infini de transformations*, Acad. Roy. Belgique. Cl. Sci. Mém. Coll. in 8° **26, fasc. 6** (1951), 1–102.

[Hum] J.E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Springer-Verlag, Berlin-Heidelberg-New York, 1972.

[Kir] A. Kirillov, *Elements of the Theory of Representations*, translated by E. Hewitt, Springer-Verlag, Berlin-Heidelberg-New York, 1976.

[Kna] A. Knapp, *Representation Theory of Real Semisimple Groups: an Overview Based on Examples*, Princeton University Press, Princeton, New Jersey, 1986.

[KV] A. Knapp and D. Vogan, *Cohomological Induction and Unitary Representations*, Princeton University Press, Princeton, New Jersey, 1995.

[Kos] 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.

[KR] B. Kostant and S. Rallis, *Orbits and representations associated with symmetric spaces*, Amer. J. Math. **93** (1971), 753–809.

[Mac] G. Mackey, *Mathematical Foundations of Quantum Mechanics*, W. A. Benjamin, Inc., New York, 1963.

[SV] S. Salamanca-Riba and D. Vogan, *On the classification of unitary representations of reductive Lie groups*, to appear in Ann. of Math.

[Tor] P. Torasso, *Quantification géométrique, opérateurs d'entrelacement et représentations unitaires de* $\widetilde{SL}_3(\mathbb{R})$, Acta Math. **150** (1983), 153–242.

[Tor2] P. Torasso, *Méthode des orbites de Kirillov-Duflo et représentations minimales des groupes simples sur un corps local de caractéristique nulle*, Duke Math. J. **90** (1997), 261–377.

[Vrg] M. Vergne, *Instantons et correspondance de Kostant-Sekiguchi*, C. R. Acad. Sci. Paris Sér. I Math. **320** (1995), 901–906.

[VG] D. Vogan, *Representations of Real Reductive Lie Groups*, Birkhäuser, Boston-Basel-Stuttgart, 1981.

[VoCI] D. Vogan, *Unitarizability of certain series of representations*, Ann. of Math. **120** (1984), 141–187.

[VO] D. Vogan, *Unitary Representations of Reductive Lie Groups*, Annals of Mathematics Studies, Princeton University Press, Princeton, New Jersey, 1987.

[VAV] D. Vogan, *Associated varieties and unipotent representations*, Harmonic Analysis on Reductive Groups (W. Barker and P. Sally, eds.), Birkhäuser, Boston-Basel-Berlin, 1991, pp. 315–388.

[VUR] D. Vogan, *Unipotent representations and cohomological induction*, The Penrose Transform and Analytic Cohomology in Representation Theory (M. Eastwood, J. Wolf, and R. Zierau, eds.), Contemporary Mathematics **154**, American Mathematical Society, Providence, Rhode Island, 1993, pp. 47–70.

[VDK] D. Vogan, *The orbit method and unitary representations of reductive Lie groups*, Algebraic and Analytic Methods in Representation Theory (Sønderborg, 1994), Perspectives in Mathematics, vol. 17, Academic Press, San Diego, 1997, pp. 243–339.

[VLC] D. Vogan, *A Langlands classification for unitary representations*, Analysis on Homogeneous spaces and Representations of Lie groups, Advanced Studies in Pure Mathematics, volume 26, Kinokuniya Company, Ltd., Tokyo, 1998.

[Wey] H. Weyl, *The Theory of Groups and Quantum Mechanics*, Dover, 1931.

[Wng] H. Wong, *Dolbeault cohomological realization of Zuckerman modules associated with finite rank representations*, J. Funct. Anal. **129** (1995), 428–454.