C^*-ALGEBRAS AND KIRILLOV'S COADJOINT ORBIT METHOD

DAVID SCHWEIN

One of the main goals of representation theory is to understand the unitary dual of a topological group, that is, the set of irreducible unitary representations. Much of modern number theory, for instance, is concerned with describing the unitary duals of various reductive groups over a local field or the adeles, and here our understanding of the representation theory is far from complete.

For a different class of groups, the nilpotent Lie groups, A. A. Kirillov gave in the mid-nineteenth century [Kir62] a simple and transparent description of the unitary dual: it is the orbit space under the coadjoint action of the Lie group on the dual of its Lie algebra. The goal of this article, notes for a talk, is to explain Kirillov's result and illustrate it with the Heisenberg group, following Kirillov's excellent and approachable book on the subject [Kir04]. We begin with an introductory section on the unitary dual of a C^*-algebra, the proper setting (currently) for unitary representations of locally compact groups, following Dixmier's exhaustive monograph on C^*-algebras [Dix77].

1. C^*-algebras and the unitary dual

The theory of unitary representations of locally compact topological groups, for instance, reductive p-adic groups, is a special case of the more general theory of representations of C^*-algebras. In this section we review the representation theory of C^*-algebras and see how it specializes to that of topological groups.

1.1. Definitions and examples.

A *Banach algebra* is a Banach space A equipped with an algebra structure with respect to which the norm is sub-multiplicative:

$$\|a \cdot b\| \leq \|a\| \cdot \|b\|, \qquad a, b \in A.$$

We do note require Banach algebras to be unital, and in fact, we will see shortly that there are many natural examples that are not unital.

Definition 1. A *C^*-algebra* is a Banach algebra A equipped with a conjugate-linear anti-involution $*$ satisfying the C^*-identity: for all $a \in A$,

$$\|a^* a\| = \|a\|^2.$$

A *morphism of C^*-algebras* is a continuous algebra map commuting with $*$.

Example 2. Let H be a Hilbert space.[1] Let $\mathrm{B}(H)$ denote the algebra of bounded operators on H, The operator norm

$$\|T\| := \sup_{\|v\|=1} \|T(v)\|$$

makes $\mathrm{B}(H)$ into a Banach space, though not a Hilbert space if $\dim H = \infty$. The algebra operations on $\mathrm{B}(H)$ are obvious, and the Hermitian transpose makes $\mathrm{B}(H)$ into a C^*-algebra.

Date: 12 February 2020.

[1] For simplicity, we assume in this article that all Hilbert spaces are separable.

The C^*-identity is surprisingly powerful. For instance, it implies that the norm is uniquely determined by the algebra structure, via the spectral radius formula

$$\|a^*a\| = \sup\{|\lambda| : a^*a - \lambda \text{ is not invertible}\}.$$

(More generally, if b is Hermitian then $\|b\|$ equals the spectral radius of b.) The identity also implies that any morphism of C^*-algebras is a contraction and that any injective morphism is an isometry.

Definition 3. Let A be a C^*-algebra. A *representation* of A is a pair (H, π) consisting of a Hilbert space H and a morphism $\pi : A \to B(H)$ of C^*-algebras.

The set $\widehat{A}$ of isomorphism classes of unitary representations of A is called the *spectrum* or *unitary dual* of A.

In this article we assume for simplicity that all representations are *nondegenerate*, meaning that no vector is annihilated by all elements of the algebra.

There is a canonical topology on $\widehat{A}$ called the *Fell topology*, named after J. M. G. Fell for a 1960 paper [Fel60] in which he studied this topology, though it appears to me that the definition is the work of several authors. Properly defining the Fell topology is beyond the scope of this article, and we have little to say about it. But from now on we will think of $\widehat{A}$ as not just a set, but a topological space.

To start with, let's try to understand what a commutative C^*-algebra looks like. The short answer is that all such algebras come from topological spaces.

Example 4. Let X be a locally compact Hausdorff space. A continuous function $f : X \to \mathbb{C}$ *vanishes at infinity* if for every $\varepsilon > 0$, there is a compact subset $K \subseteq X$ such that $\sup_{x \in X \setminus K} |f(x)| < \varepsilon$. Let $C_0(X)$ denote the space of complex functions on X that vanish at infinity. It is a C^*-algebra with L^∞-norm and complex conjugation of functions. The algebra $C_0(X)$ is unital if and only if X is compact. Every point $x \in X$ gives rise to a character $\mathrm{ev}_x : C_0(X) \to \mathbb{C}$, evaluation at x. The assignment $x \mapsto \mathrm{ev}_x$ is a continuous map $X \to \widehat{C_0(X)}$, which turns out to be a homeomorphism.

In the opposite direction, starting from a commutative C^*-algebra A, there is a natural evaluation map $A \to C_0(\widehat{A})$ sending f to the function $\xi \mapsto \xi(f)$.

Theorem 5 (Gelfand)**.** *For every commutative C^*-algebra A, the evaluation map $A \to C_0(\widehat{A})$ is an isomorphism of C^*-algebras. Moreover, the functors $X \mapsto C_0(X)$ and $A \mapsto \widehat{A}$ define an equivalence between the category of compact Hausdorff spaces and the category of commutative unital C^*-algebras with unital morphisms.*

Gelfand's theorem is extraordinarily important. For one, it inspired the definition of the prime spectrum in algebraic geometry, formulated by Grothendieck and collaborators. For another, since the theorem implies that locally compact Hausdorff spaces are more or less equivalent to a special class of C^*-algebras, it suggests thinking of noncommutative C^*-algebras as noncommutative topological spaces. This perspective motivated Alain Connes's profound theory of noncommutative geometry.

Remark 6. The theorem extends to an equivalence between the categories of locally compact Hausdorff spaces with proper maps and the category of commutative C^*-algebras with nondegenerate morphisms. Here a morphism is *nondegenerate* if it sends approximate identities to approximate identities. The basic example this property is meant to exclude is the zero map $C_0(Y) \to C_0(X)$, which does not correspond to any map of topological spaces $X \to Y$.

Alternatively, the theorem generalizes to an equivalence between the categories of pointed compact Hausdorff spaces and the category of augmented C^*-algebras. The connection between the two generalizations is the one-point compactification functor on the topological side and the unitalization functor on the C^*-algebra side.

Example 7. There is a construction, which we will discuss soon in more detail, that assigns to any locally compact topological group G a C^*-algebra $C^*(G)$. This construction has the property that the representations of G are the same as the representations of $C^*(G)$. When G is a finite group, however, there are no complications and $C^*(G) = \mathbb{C}[G]$, the usual complex group algebra of G.

More generally, given a locally compact topological group G acting on a locally compact Hausdorff space X, one can form a C^*-algebra $C^*(G \curvearrowright X)$ which morally represents the functions on the "noncommutative space" $G\backslash X$. This C^*-algebra carries much more information than the topological quotient $G\backslash X$, which is often non-Hausdorff. More generally, one can form a C^*-algebra from any groupoid [Wil19].

1.2. **Representations of locally compact groups.** In this subsection we specialize the theory of C^*-algebras to study the (unitary) representation theory of locally compact groups G. For clarity of exposition we assume that G is unimodular. Given a Hilbert space H, let $U(H) \subseteq \mathrm{Aut}(H)$ denote the group of (bounded) unitary linear operators on H.

Definition 8. A (unitary) representation (π, H) of G is a Hilbert space H together with a homomorphism $\pi : G \to U(H)$ such that the induced map $G \times H \to H$ is continuous.

To apply the theory of C^*-algebras to G, we need to construct a C^*-algebra out of it. The natural candidate is a space of functions on G where multiplication is given by convolution. In order for this space of functions to be a Banach algebra, we must equip it with a norm that behaves well with respect to convolution. The simplest possibility is an L^p-norm. Moreover, the generalized Young's inequality

$$\|f_1\|_p \cdot \|f_2\|_q \leq \|f_1 * f_2\|_r, \qquad \frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r},$$

shows that the space $L^1(G)$ is closed under convolution. We further endow this space with the anti-involution

$$f^*(x) := \bar{f}(x^{-1}).$$

Our first candidate for $C^*(G)$ is therefore $L^1(G)$. This is a very good candidate because its representation theory reflects very well the representation theory of G. The only problem is that $L^1(G)$ is not a C^*-algebra due to the failure of the C^*-identity, but this turns out to be not so great of a difficulty.

Remark 9. To avoid the choice of a Haar measure, one can replace $L^1(G)$ with a certain space of complex measures on G. This is surely the better choice, but we have avoided discussing it in detail for simplicity.

Theorem 10. *There is a canonical bijection between representations of G and representations of $L^1(G)$.*

Proof sketch. The rough idea is that the representation (π, H) of G should correspond to the representation of $L^1(G)$ on H for which the function f acts on $v \in H$ by

$$\pi(f)(v) = \int_G f(x)\pi(x)(v)\,\mathrm{d}x.$$

This is only a rough idea because the integrand takes values in a Hilbert space, so the integral must be interpreted carefully. □

Example 11. The group G acts on $L^2(G)$ by multiplication:

$$\rho(g)(v) : x \mapsto v(g^{-1}x), \qquad v \in L^2(G).$$

The extension of this action to $L^1(G)$ is convolution, which is defined by Young's inequality:

$$\rho(f)(v) = f * v.$$

Let's give an example of this correspondence, while simultaneously explaining (abstractly) the failure of $L^1(G)$ to be a C^*-algebra.

Example 12. Suppose G is abelian and consider the representation $(\rho, L^2(G))$ of the previous example. We can use the Fourier transform to describe it in a different way. Recall that for $f \in L^1(G)$, the Fourier transform $\hat{f} : G \to \mathbb{C}$ is defined by

$$\hat{f}(\xi) := \xi(f).$$

The theory of Pontryagin duality shows that the Fourier transform gives an isomorphism between $L^2(G)$ and $L^2(\widehat{G})$, and that this isomorphism is G-equivariant if we endow $L^2(\widehat{G})$ with the action

$$\hat{\rho}(f)(v) := \hat{f}v.$$

It follows that $\|\rho(f)\| = \|\hat{\rho}(f)\| = \|\hat{f}\|_\infty$.

We can now show abstractly, by contradiction, that $L^1(G)$ does not satisfy the C^*-identity. Consider the representation as a morphism $\rho : L^1(G) \to \mathrm{B}(L^2(G))$. It can be shown that ρ is injective, in other words, faithful. So if $L^1(G)$ were a C^*-algebra then ρ, like every injective morphism of C^*-algebras, would be an isometry. But generally $\|\hat{f}\|_\infty \neq \|f\|_1$.

Even though $L^1(G)$ is not a C^*-algebra, there is a general and universal construction that transforms an algebra like $L^1(G)$ (specifically, a Banach $*$-algebra that admits an approximation to the identity) into a C^*-algebra. The previous example suggests that we should modify the norm on $L^1(G)$ by taking into account the norms in the irreducible representations of G. If the new norm $\|\cdot\|_{C^*}$ satisfies the C^*-identity then we can complete with respect to the norm to produce a C^*-algebra. So define the norm $\|\cdot\|_{C^*}$ on $L^1(G)$ by

$$\|f\|_{C^*} := \sup_{\pi \in \widehat{G}} \|\pi(f)\|,$$

as in the previous example, and let $C^*(G)$ be the completion of $L^1(G)$ with respect to $\|\cdot\|_{C^*}$. General properties of this construction ensure that it has no effect on the representation theory. That is, there is a canonical equivalence of categories between the representations of $L^1(G)$ and the representations of $C^*(G)$.

Example 13. Let's see how the construction of $C^*(G)$ works for $G = \mathbb{Z}$. Since the dual group of $\mathbb{Z}$ is the circle $S^1 = \{c \in \mathbb{C} : |c| = 1\}$, we expect that $C^*(\mathbb{Z}) = C(S^1)$.

The space $L^1(\mathbb{Z})$ can be identified with the algebra of Laurent series $f(z) = \sum_{n\in\mathbb{Z}} a_n z^n$ such that $\sum_n |a_n| < \infty$. In this notation, the $*$-operator acts by

$$f^*(z) = \bar{f}(z^{-1}).$$

In this description, it is easy to cook up examples showing that $L^1(\mathbb{Z})$ is not a C^*-algebra: for instance, the function $f(z) = 1 + i(z + z^{-1})$ has $\|f^2\| = 7$ while $\|f^*f\| = 3$.[2] Every Laurent series represents a continuous function on the circle, giving a natural inclusion $L^1(\mathbb{Z}) \subset C(S^1)$. However, this inclusion is *not* an equality: there are examples of continuous periodic functions whose Fourier coefficients do not lie in $L^1(\mathbb{Z})$ [Kat04, 2.4]. Apparently there is no simple characterization of the Fourier coefficients of continuous functions.

Let χ_c denote the character of $\mathbb{Z}$ corresponding to $c \in S^1$. Then $\chi_c f = f(c)$, so that $\|f\|_{C^*} = \|f\|_{L^\infty(S^1)}$. That is, the C^*-norm on $L^1(\mathbb{Z})$ is just the sup norm. Since $L^1(\mathbb{Z})$, thought of a space of functions on S^1, contains all polynomial functions, and these are dense in $C(S^1)$, the completion of $L^1(\mathbb{Z})$ with respect to this norm is $C(S^1)$, as expected.

1.3. **Plancherel measure.** In the theory of Pontryagin duality, where the group G is abelian, a choice of Haar measure on G yields a canonical measure on $\widehat{G}$. This subsection explains how the unitary dual admits a similar measure, the Plancherel measure, when G is nonabelian.

Before discussing the Plancherel measure, let's consider the more general problem of imposing a measure on the unitary dual of a C^*-algebra, and the relationship of the measure to representation theory.

Remark 14. The theory that follows works best when the C^*-algebra is *postliminal*, a technical condition that we will not define, but that includes commutative C^*-algebras and the C^*-algebras of the groups of interest to us later. So we assume from now on that the C^*-algebras under consideration are postliminal, and that the groups under consideration have postliminal C^*-algebras. This assumption does not include every locally compact group G, but it does include nilpotent Lie groups [Kir62, Theorem 7.3], reductive Lie groups [HC53, Theorem 7], and reductive p-adic groups [HC70, Theorem 4].

The key technical tool in this area is the direct integral, whose rough idea is as follows. Suppose we have a measurable space X and for each $x \in X$, a Hilbert space H_x. For each measure μ on X, we can construct the *direct integral*

$$\int_X^{\oplus} H_x \,\mathrm{d}\mu(x).$$

It is defined as the set of measurable sections $v = (v_x)_{x\in X} \in \prod_{x\in X} H_x$ that are *square-integrable* with respect to μ, meaning that

$$\|v\|_\mu^2 := \int_X \|v_x\|^2 \,\mathrm{d}\mu(x) < \infty.$$

If $f : X \to \mathbb{R}_{\geq 0}$ is a nonnegative function whose vanishing set has measure zero then the direct integral will be the same whether we form it with respect to the measure μ or the measure $f\mu$. So the direct integral depends only weakly on the measure.

[2]We can also compute that $\|\hat{f}\|_\infty = \sqrt{5} < 3 = \|f\|_1$.

Example 15. Let X be a locally compact Hausdorff space with C^*-algebra $C_0(X)$ and let μ be a Borel measure on X. Recall that X is isomorphic to the spectrum of $C_0(X)$, the point $x \in X$ corresponding to the evaluation character $\mathrm{ev}_x : C_0(X) \to \mathbb{C}$ sending f to $f(x)$. Then

$$\int_X^{\oplus} \mathrm{ev}_x \, \mathrm{d}\mu(x) = L^2(X, \mu).$$

This example suggests that the direct integral of irreducible representations of a noncommutative C^*-algebra can be interpreted as the L^2-space of the noncommutative space $\widehat{A}$. How, then, could we interpret the L^p-space of $\widehat{A}$ for $p \neq 2$?

Here is the general theorem on measures.

Theorem 16 ([Dix77, 8.6.6])**.** *Let A be a postliminal C^*-algebra. For each representation π of A, there are mutually singular measures $\mu_1, \mu_2, \ldots, \mu_\infty$ on $\widehat{A}$ such that*

$$\pi \simeq \int_{\widehat{A}} \xi \, \mathrm{d}\mu_1(\xi) \oplus 2 \int_{\widehat{A}} \xi \, \mathrm{d}\mu_2(\xi) \oplus \cdots \oplus \mathbb{N} \int_{\widehat{A}} \xi \, \mathrm{d}\mu_\infty(\xi).$$

The measures μ_i in the theorem are not uniquely determined by A, but they are uniquely determined up to equivalence of measures: for each i, the collection of subsets of $\widehat{A}$ that have μ_i-measure zero is uniquely determined. Moreover, giving each of the measures μ_i is the same as giving their sum $\mu := \mu_1 + \mu_2 + \cdots + \mu_\infty$.

To summarize, every representation of a postliminal C^*-algebra A gives rise to a class of measures on the spectrum $\widehat{A}$, and conversely, every measure on $\widehat{A}$ gives rise to a representation. So there is a weak correspondence between representations of A and measures on $\widehat{A}$.

The rough idea of the Plancherel measure is that it should decompose the regular representation $L^2(G)$. Although this condition turn out to be true, it does not uniquely pin down the measure, since the disintegration theory of representations yields a huge class of measures on $\widehat{G}$ corresponding to a representation

$$\int_{\widehat{G}}^{\oplus} \xi \, \mathrm{d}\mu(\xi).$$

To circumvent this difficulty, we use the Plancherel formula.

Recall that when G is abelian, the Plancherel formula states that the Fourier transform is an L^2-isometry: for any $f \in L^1(G) \cap L^2(G)$,

$$\int_G |f(x)|^2 \, \mathrm{d}x = \int_{\widehat{G}} |\hat{f}(\xi)|^2 \, \mathrm{d}\xi,$$

where $\hat{f}$ is the Fourier transform of f. To generalize the formula to noncommutative G, we need some notion of the Fourier transform.

Definition 17. The *Fourier transform* $\hat{f}$ of $f \in L^1(G)$ is the "function" $\hat{f}$ defined by

$$\hat{f}(\pi) := \pi(f) = \int_G f(x)\pi(x) \, \mathrm{d}x \in B(H_\pi).$$

This is the same formula as in the commutative case, but now $\hat{f}$ is not a function in the usual sense because the space in which the operator $f(\pi)$ lives depends on π.

Theorem 18. *For each Haar measure μ on G there exists a unique measure $\hat{\mu}$ on $\widehat{G}$, the* Plancherel measure, *such that for all $f \in L^1(G) \cap L^2(G)$,*

$$\int_G |f(x)|^2 \, \mathrm{d}\mu(x) = \int_{\widehat{G}} \mathrm{Tr}(\hat{f}(\xi)\hat{f}(\xi)^*) \, \mathrm{d}\hat{\mu}(\xi).$$

Moreover, the Plancherel measure disintegrates $L^2(G)$:

$$L^2(G) = \int_{\widehat{G}}^{\oplus} \xi \otimes \xi^* \, \mathrm{d}\hat{\mu}(\xi).$$

In general, the Plancherel measure on $\widehat{G}$ may not be supported on all of $\widehat{G}$. For example, the unitary dual of $\mathrm{SL}_2(\mathbb{R})$ contains an interval of irreducible representations, the *complementary series*, whose measure is zero. The support of the Plancherel measure is called the *reduced dual* $\widehat{G}_r$. Typically the reduced dual is much easier to understand than the entire unitary dual. For one, a modified construction the C^*-algebra yields a C^*-algebra $C^*_r(G)$, the *reduced C^*-algebra* of G, whose spectrum is $\widehat{G}_r$: we define $C^*_r(G)$ as the completion of $L^1(G)$ with respect to the norm

$$\|f\|_{C^*_r} := \|\rho(f)\|$$

where ρ is the regular representation of G on $L^2(G)$.

2. Coadjoint orbit method

In this section we describe the coadjoint orbit method, paying special attention to the case of the Heisenberg group. Although the method is a general philosophy that applies to many classes of groups, we restrict attention to the (simply-connected) nilpotent Lie groups, a class for which the method works particularly well.

2.1. **Motivation: p-groups.** Before thinking about Lie groups, we consider a simpler toy model in finite groups that illustrates many of the phenomena to come. Recall that a finite group is a *p-group* if its order is a power of p. The following example is most relevant to us.

Example 19. Let $\underline{U}$ be a unipotent algebraic group over the finite field $\mathbb{F}_q$ of characteristic p. Then $\underline{U}(\mathbb{F}_q)$ is a p-group.

The structure theory of p-groups is extraordinarily complicated. Its complexity has obstructed all attempts to classify such groups or their unitary duals. We can say something weaker, though. Recall that a representation is *monomial* if it is induced from a character.

Theorem 20. *Every irreducible representation of a p-group is monomial.*

The key property used in the proof is that the center of a p-group is nontrivial. We can therefore leverage the center to prove the theorem by induction on the cardinality of the group.

Proof. Let G be the p-group and (π, V) an irreducible representation of it. When G is abelian the result is clear, so without loss of generality we may take π to be faithful and G to be nonabelian. Let A be a normal abelian subgroup of G which is not contained in the center. (That such a subgroup exists follows from the fact that all groups of order p^2 are abelian.) Consider the isotypic decomposition of ρ restricted to A:

$$V = \bigoplus_{i \in I} V_i.$$

Since the representation is faithful, there must be at least two isotypic components: otherwise A would act as scalars on V, hence centralize $\mathrm{GL}(V)$, hence centralize G. Choose some isotypic component V_0 and let H be the stabilizer of V_0, so that π is induced from the H-representation $(\pi|_H, V_0)$. We are now finished by induction because H is a proper subgroup of G, as $|I| > 1$. □

Remark 21. The conclusion of the theorem holds more generally for supersolvable groups, the smallest nontrivial class of finite groups closed under cyclic extensions [Ser77, Section 8.5].

To describe the unitary dual further, one would need a good handle on the subgroups of G, and this is generally possible only in special cases, since the structure theory of p-groups is so complicated. The goal of this section is to describe an analogue of this theory for connected Lie groups. In that setting, we can use the exponential map to describe the subgroups, and in addition to showing that all representations are monomial, we can describe the unitary dual explicitly.

2.2. **Nilpotent Lie groups.** The class of nilpotent Lie groups is the smallest class of Lie groups that contains real vector spaces and is closed under connected central extensions.

Definition 22. A Lie group G is *nilpotent* if it admits a filtration

$$1 = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_n = G$$

of connected normal subgroups such that for all i, the quotient G_{i+1}/G_i is central in G/G_i.

Example 23. Let $\underline{U}$ be a unipotent algebraic group over $\mathbb{R}$. Then $\underline{U}(\mathbb{R})$ is a nilpotent simply-connected Lie group. Of particular importance is the unipotent group U_n of upper triangular matrices. One can use Ado's theorem to show that any simply-connected nilpotent Lie group embed in U_n. (This does not seem to immediately imply that every such Lie group is the group of $\mathbb{R}$-points of a unipotent group, and I don't know if this stronger statement is true.)

Not all nilpotent groups are simply connected: the easiest example is the circle group $\mathbb{R}/\mathbb{Z}$.

Our understanding of the representation theory of nilpotent Lie groups traces to the exponential map.

Theorem 24. *Let G be a simply-connected nilpotent Lie group. Then the exponential map* $\exp : \mathfrak{g} \to G$ *is a diffeomorphism establishing a bijection between sub-Lie algebras* $\mathfrak{h} \subseteq \mathfrak{g}$ *and closed connected subgroups* $H \subseteq G$.

In order to use the theorem we must therefore restrict our attention to simply-connected nilpotent groups, a class of group works which the orbit method works perfectly.

For this article our motivating important example of a nilpotent Lie group is the Heisenberg group, whose definition we now recall.

A *symplectic space* is a (real, finite-dimensional) vector space V endowed with a nondegenerate symplectic form. Symplectic spaces are classified by their dimension, an even number. In general, the (central) extensions of a group G by an abelian group A are classified by the two-cocycles $G \times G \to A$, taken up to coboundaries. It turns out that every symplectic space gives rise to such a cocycle, namely, the symplectic form: bilinearity implies that the form is a cocycle. This construction produces the Heisenberg group, though for normalization reasons, we halve the form.

Definition 25. The *Heisenberg group* $H(V)$ of the symplectic space $(V, \langle -,-\rangle)$ is the central extension $1 \to \mathbb{R} \to H(V) \to V \to 1$ classified by the symplectic form $\frac{1}{2}\langle -,-\rangle : V \times V \to \mathbb{R}$.

In other words, the Heisenberg group $H(V)$ has underlying set $V \times \mathbb{R}$ with multiplication

$$(v,t) \cdot (v',t') = (v+v', t+t'+\tfrac{1}{2}\langle v, v'\rangle).$$

The center of the Heisenberg group is the factor $\mathbb{R}$: hence $H(V)$ is nilpotent.

Since $H(V)$ is nilpotent, we should be able to express its elements in coordinates as upper-triangular matrices. One way to do this is as follows. A *polarization* of the symplectic space V is an orthogonal decomposition $V = W \oplus W'$ such that $W^\perp = W$ and $(W')^\perp = W'$. (There is no canonical polarization: indeed, the space of polarizations is in bijection with a certain conjugacy class of parabolic subgroups of $\mathrm{Sp}(V)$.) The symplectic form identifies W' with the dual space of W, and we can use the polarization to interpret elements of $H(V)$ as 3×3 block matrices

$$\begin{bmatrix} 1 & w & u \\ & 1 & w^\vee \\ & & 1 \end{bmatrix}, \qquad w \in W, \quad w^\vee \in W^\vee, \quad t \in \mathbb{R}.$$

If you like, the group can be fully coordinatized by further choosing a basis for W, which then gives a basis for W'. The two descriptions of $H(V)$ are related by the equations

$$v = w + w^\vee, \qquad t = u - \tfrac{1}{2}\langle w, w^\vee\rangle.$$

The Lie algebra $\mathfrak{h}(V)$ of the Heisenberg group $H(V)$ has underlying vector space $V \oplus \mathbb{R}$ with bracket

$$[(x,s),(x',s')] = (0, \langle x, x'\rangle).$$

(This formula explains the appearance of the factor $\frac{1}{2}$ in our formula for multiplication in the Heisenberg group; if we had omitted the $\frac{1}{2}$ there, then the symplectic form here would have to be doubled to compensate.) In this way we can recover the symplectic structure on V from the group structure on $H(V)$. The coordinate description of $H(V)$ shows that the exponential map $\exp : \mathfrak{h}(V) \to H(V)$ is given by the simple formula

$$\exp(x,s) = (x,s).$$

2.3. **Coadjoint orbits.** Our next goal is to define the space of coadjoint orbits of a nilpotent Lie group G, illustrating the construction with the Heisenberg group. The key construction is the *coadjoint representation* of G, the action of G on its dual Lie algebra $\mathfrak{g}^\vee$.

Example 26. When $G = \underline{G}(\mathbb{R})$ is the group of $\mathbb{R}$-points of a reductive algebraic group $\underline{G}$, the adjoint representation $\mathfrak{g}$ is self-dual. When G is semisimple duality follows from the observation that the representation preserves the Killing form, and in general, the adjoint representation factors through the adjoint quotient of G.

In particular, the coadjoint representation of $\mathrm{GL}(V)$ is just the vector space $\mathrm{End}(V)$ with conjugation action. The coadjoint orbit space $\mathrm{End}(V)/\mathrm{GL}(V)$ can be described explicitly using eigenvalues and (some weak generalization of) the Jordan normal form. We can also use the example of $\mathrm{GL}(V)$ to compute the coadjoint representation of nilpotent Lie groups by embedding them into $\mathrm{GL}(V)$, with the caveat that the adjoint representation of such a group may not be self-dual.

For the Heisenberg group, the adjoint action of $H(V)$ on $\mathfrak{h}(V)$ is given by the formula

$$\mathrm{Ad}(v,t)(x,s) = (x, \langle v, x\rangle + s).$$

The symplectic form on V gives a canonical identification of V with its dual $V^\vee$, hence an identification of $\mathfrak{h}(V)$ with $\mathfrak{h}(V)^\vee$, provided we make the standard identification $\mathbb{R} = \mathbb{R}^\vee$ via the multiplication pairing. With respect to this pairing, the coadjoint action must satisfy

$$\langle \mathrm{Ad}^\vee(v,t)(x,s), (x',s')\rangle = \langle (x,s), \mathrm{Ad}(-v,-t)(x',s')\rangle.$$

The righthand side of this expression is

$$\langle x, x'\rangle + s(s' - \langle v, x'\rangle) = \langle x - sv, x'\rangle + ss'$$

Hence the coadjoint action is given by the formula

$$\mathrm{Ad}^\vee(v,t)(x,s) = (x - sv, s).$$

Unlike the general linear group, the Heisenberg group's adjoint representation is not self-dual.

We can now describe the space $H(V)\backslash\mathfrak{h}(V)^\vee$ of coadjoint orbits. There are two cases. If $(x,s) \in \mathfrak{h}(V)^\vee$ has $s = 0$ then it lies in its own orbit. Otherwise, if $s \neq 0$, all elements of the form (x,s), for varying $x \in V$, lie in the same orbit, the affine hyperplane (V,s). It follows that as a set, the orbit space is the disjoint union $H(V)\backslash\mathfrak{h}(V)^\vee = V \sqcup \mathbb{R}^\times$, depicted at right. The topology on the quotient is not the coproduct topology however: the quotient is non-Hausdorff because any sequence of elements in the copy of the reals that converges to zero converges to every other element of $V^\vee$ as well.

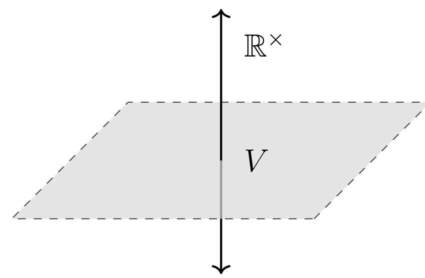

In the case of the Heisenberg group, the coadjoint have dimension either zero or $\dim V$, and in particular are even. Further, the orbits carry a canonical symplectic structure. This is a general feature of coadjoint orbits of nilpotent groups.

Theorem 27. *Let G be a simply-connected nilpotent group. and let Ω be a coadjoint orbit of G. Then Ω carries a canonical G-invariant symplectic form. In particular, $\dim\Omega$ is even.*

Proof sketch. Since the symplectic form we are searching for is G-invariant, it is enough to define it at a single point $x^\vee \in \Omega$. Let H be the stabilizer of $x^\vee$ in G and let $\mathfrak{h}$ be its Lie algebra. The choice of basepoint $x^\vee$ yields an identification $T_{x^\vee}(\Omega) \simeq \mathfrak{g}/\mathfrak{h}$. Now define the bilinear form $B_{x^\vee}$ on $\mathfrak{g}$ by

$$B_{x^\vee}(x,y) = \langle x^\vee, [x,y]\rangle.$$

It is an easy exercise to check that (1) the form $B_{x^\vee}$ is alternating, (2) the kernel of $B_{x^\vee}$ is $\mathfrak{h}$, and (3) the form $B_{x^\vee}$ is H-invariant. Hence $B_{x^\vee}$ defines a nondegenerate 2-form.

These conditions are already enough to imply that $\dim\Omega$ is even. However, a symplectic form is required to satisfy an additional condition, namely, that it be a *closed* 2-form. Checking this additional condition requires more work, and one needs it to access tools from symplectic topology. □

2.4. **Unitary dual: Heisenberg group.** It turns out that for nilpotent groups, the construction of the previous subsection exactly describes the unitary dual.

Theorem 28. *Let G be a simply-connected nilpotent Lie group. Then there is a canonical homeomorphism*

$$G\backslash\mathfrak{g}^\vee \to \widehat{G}, \qquad \Omega \mapsto \pi_\Omega.$$

As we will see, the homeomorphism preserves much more structure than the topology.

Remark 29. The description of the unitary dual in terms of the space of coadjoint orbits is visibly false when the group is not simply connected, as we can see already for the nilpotent group $S^1 = \mathbb{R}/\mathbb{Z}$: its unitary dual is $\mathbb{Z}$ but its coadjoint orbit space is $\mathbb{R}$.

Before explaining the bijection between the unitary dual and the space of coadjoint orbits, let's see what happens for the Heisenberg group. To classify the irreducible unitary representations (π, H) of $H(V)$ we examine their central character, recalling that the center of $H(V)$ is $\mathbb{R}$.

There is a dichotomy here: either the central character is trivial or it is nontrivial. If the central character is trivial then π is inflated from a representation of the quotient $H(V)/\mathbb{R} = V$; these are classified by Pontryagin duality, and are all one-dimensional. If the central character is nontrivial then the Stone–von Neumann theorem states that it uniquely determines the representation up to isomorphism.

Theorem 30 (Stone–von Neumann)**.** *The assignment $(\pi, H) \mapsto \pi|_{\mathbb{R}}$ yields a bijection between the nontrivial characters of $\mathbb{R}$ and the irreducible unitary representations of $H(V)$ with nontrivial central character.*

This analysis gives the bijection $\widehat{H(V)} \simeq H(V)\backslash\mathfrak{h}(V)^\vee = V \sqcup \mathbb{R}^\times$: the elements of V correspond to the representations inflated from characters of V, and the elements of $\mathbb{R}$ correspond to the Stone–von Neumann representations.

The construction of the Stone–von Neumann representations illustrates the general description of the unitary dual in terms of the coadjoint space. Given a real number $a \in \mathbb{R}$, the corresponding Stone–von Neumann representation π_a depends on a choice of polarization $V = W \oplus W'$, though any two choices yield isomorphic representations. Define the representation $(\pi_a, \mathrm{L}^2(W))$ of $H(V)$ by

$$(\pi_a(w + w', t)f)(x) = e(at + \langle x, w'\rangle)f(x + w), \qquad w \in W, w' \in W'$$

where $e(t) := \exp(2\pi i t)$ for $t \in \mathbb{R}$. More intrinsically, the representation $(\pi_a, \mathrm{L}^2(W))$ can be described as the representation induced from the inflation of $t \mapsto e(at)$ to the subgroup $W \times \mathbb{R}$ of $H(V)$. The fact that $W^\perp = W$ implies that $W \times \mathbb{R}$ is abelian.

2.5. **Unitary dual: nilpotent groups.** How does the bijection of the Stone–von Neumann theorem generalize to other nilpotent groups? We start by describing the characters.

Lemma 31. *Every character of a simply-connected nilpotent Lie group G has the form*

$$\rho_{x^\vee}(\exp x) = e(\langle x^\vee, x\rangle)$$

for some linear function $x^\vee \in \mathfrak{g}^\vee$ such that $[\mathfrak{g}, \mathfrak{g}] \subseteq \ker x^\vee$.

Proof. Since G is simply-connected, any character $G \to S^1$ lifts to a homomorphism $G \to \mathbb{R}$. And because of the tight connection between nilpotent Lie groups and their Lie algebras, via the exponential map, giving such a homomorphism amounts to giving a Lie algebra morphism $\mathfrak{g} \to \mathbb{R}$. □

The characters of the lemma are precisely the characters whose inductions fill out the unitary dual. But in order for the induced representation to be irreducible, we must impose a maximality condition.

Definition 32. Let $\mathfrak{g}$ be a Lie algebra, $\mathfrak{h} \subseteq \mathfrak{g}$ a subalgebra, and $x^\vee \in \mathfrak{g}^\vee$ a functional. The subalgebra $\mathfrak{h}$ is *subordinate to* $x^\vee$ if $[\mathfrak{h}, \mathfrak{h}] \subseteq \ker x^\vee$, and is *maximally subordinate to* $x^\vee$ if, in addition, it has maximal dimension among all such subalgebras.

Theorem 33. *Let G be a simply-connected nilpotent Lie group.*

(1) $\operatorname{Ind}_H^G \rho_{x^\vee,H}$ *is irreducible if and only if $\mathfrak{h}$ is maximally subordinate to $x^\vee$.*
(2) *Every irreducible representation is of the form $\operatorname{Ind}_H^G \rho_{x^\vee,H}$ for some H and $x^\vee$.*
(3) *Two representations $\operatorname{Ind}_H^G \rho_{x^\vee,H}$ and $\operatorname{Ind}_K^G \rho_{y^\vee,K}$ are isomorphic if and only if $x^\vee$ and $y^\vee$ lie in the same coadjoint orbit.*

Write π_Ω for the irreducible representation of G corresponding to the coadjoint orbit Ω.

Proof sketch. The proof proceeds by induction on the dimension of G. Without loss of generality we may assume the center Z of G has dimension one. Kirillov uses as a crucial technical tool a particular structure in G which he calls a *canonical decomposition*: a decomposition

$$\mathfrak{g} = \mathbb{R}x + \mathbb{R}y + \mathbb{R}z + W, \qquad x, y \in \mathfrak{g}, \quad z \in \mathfrak{z}, \quad W \subset \mathfrak{g}.$$

such that $[x, y] = z$ and $[y, W] = 0$. In spite of the name, such a decomposition is far from canonical. It is not hard to see that canonical decompositions always exist. The span of $\{x, y, z\}$ is a three-dimensional Heisenberg Lie algebra, and Kirillov uses its integrated Heisenberg group to study the representation theory of G. $\square$

It is a priori very difficult to compute maximal subordinate subalgebras: how does one know when the maximum dimension has been attained? Fortunately, there is a simple formula for the dimension.

Theorem 34. *Let Ω be a coadjoint orbit of a simply-connected nilpotent group G, let $x^\vee \in \Omega$, and let $\mathfrak{h} \subseteq \mathfrak{g}$ be a subalgebra maximally subordinate to $x^\vee$. Then*

$$\dim \mathfrak{h} = \dim G - \tfrac{1}{2} \dim \Omega.$$

Example 35. Returning to the Heisenberg group, let Ω be a coadjoint orbit isomorphic to $V^\vee$. Then $\mathfrak{h} = W \oplus \mathbb{R}$ where $V = W \oplus W^\vee$ is a polarization, and the equation above becomes $(\dim W + 1) = (\dim V + 1) - \frac{1}{2}(\dim V)$.

Remark 36. There is a strong formal analogy between Kirillov's description of the unitary dual of a nilpotent Lie group N and Harish-Chandra's description of the tempered dual of a reductive group G. In both cases, the duals are constructed from smaller subgroups by inflating, then inducting. The inflation-induction procedure relies on an auxiliary choice that is ultimately irrelevant up to isomorphism: for N, a maximally subordinate subalgebra, and for G, a parabolic subgroup. It would be very interesting to explore this analogy further.

2.6. **Further results.** We claimed earlier that the homeomorphism between the coadjoint orbit space and the unitary dual preserves much more structure than just the topology. In this subsection we justify this claim by explaining several additional structures preserved by the correspondence. The simplest condition is for duality: it is easy to see that $\pi_\Omega^\vee \simeq \pi_{-\Omega}$.

Next is the Plancherel measure. Since the exponential map gives a homeomorphism between $\mathfrak{g}$ and G, to describe the Haar measure on G we can describe the Haar measure on $\mathfrak{g}$ in local coordinates. It turns out that this measure is the usual Lebesgue measure on $\mathfrak{g}$: in other words, the exponential map is compatible with Haar measures. So fix a Haar measure μ on $\mathfrak{g}$ with dual measure $\mu^\vee$ on $\mathfrak{g}^\vee$. Then the Plancherel measure on $\widehat{G}$ with respect to $\exp \mu$ is exactly the pushforward of $\mu^\vee$ along the projection $\mathfrak{g}^\vee \to G\backslash\mathfrak{g}^\vee$.

For restriction and induction, let $H \subseteq G$ be a nilpotent subgroup of G. There is a natural equivariant restriction map $p : \mathfrak{g}^\vee \to \mathfrak{h}^\vee$. For restriction, if Ω is a coadjoint orbit of G then

$$\operatorname{Res}_H^G \pi_\Omega = \int_{\omega \subseteq p(\Omega)}^{\oplus} m(\omega, \Omega) \pi_\omega \, \mathrm{d}\omega;$$

for induction, if ω is a coadjoint orbit of H then

$$\operatorname{Ind}_H^G \pi_\omega = \int_{\Omega \subseteq p^{-1}(\omega)}^{\oplus} m(\omega, \Omega) \pi_\Omega \, \mathrm{d}\omega;$$

I will not describe the multiplicity function $m(\omega, \Omega)$ explicitly, except to say that it takes values in $\{1, \infty\}$ and is the same for both restriction and induction.

As for the tensor product, let Ω and Ω' be coadjoint orbits of G. Their arithmetic sum is defined to be $\Omega + \Omega' = \{a + a' : a \in \Omega, a' \in \Omega'\}$. Then

$$\pi_{\Omega+\Omega'} = \int_{\omega \subseteq \Omega + \Omega'}^{\oplus} m(\omega, \Omega, \Omega') \pi_\omega \, \mathrm{d}\omega,$$

where $m(\omega, \Omega, \Omega') \in \{1, \infty\}$ is a certain multiplicity that I won't specify. In fact, this result can be deduced from the decomposition for restriction because

$$\pi_\Omega \otimes \pi_{\Omega'} = \operatorname{Res}_G^{G \times G}(\pi_\Omega \boxtimes \pi_{\Omega'}).$$

Example 37. Let's return to the case of the Heisenberg group $H(V)$. For $a \in \mathbb{R}^\times$, let $(\pi_a, \mathrm{L}^2(W))$ denote the Stone–von Neumann representation for a, where $V = W \oplus W^\vee$ is a polarization. Consider the subgroup W of $H(V)$. The coadjoint orbit corresponding to π_a is the affine space $a + V^\vee$. It surjects onto $W^\vee$ under the restriction map. Then the theory of the Fourier transform gives the Hilbert space decomposition

$$\operatorname{Res}_W^{H(V)} \pi_a = \int_{W^\vee}^{\oplus} \chi_{w^\vee} \, \mathrm{d}w^\vee,$$

where $\chi_{w^\vee}$ is the unitary character of W corresponding to $w^\vee \in W^\vee$.

References

[Dix77] Jacques Dixmier, *C^*-algebras*, North-Holland Publishing Co., Amsterdam-New York-Oxford, 1977, Translated from the French by Francis Jellett, North-Holland Mathematical Library, Vol. 15. MR 0458185

[Fel60] J. M. G. Fell, *The dual spaces of C^*-algebras*, Transactions of the American Mathematical Society **94** (1960), 365–403. MR 146681

[HC53] Harish-Chandra, *Representations of a semisimple Lie group on a Banach space. I*, Transactions of the American Mathematical Society **75** (1953), 185–243. MR 56610

[HC70] ———, *Harmonic analysis on reductive p-adic groups*, Lecture Notes in Mathematics, Vol. 162, Springer-Verlag, Berlin-New York, 1970, Notes by G. van Dijk. MR 0414797

[Kat04] Yitzhak Katznelson, *An introduction to harmonic analysis*, third ed., Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2004. MR 2039503

[Kir62] A. A. Kirillov, *Unitary representations of nilpotent Lie groups*, Russian Mathematical Surveys **17** (1962), no. 4, 53–104.

[Kir04] ———, *Lectures on the orbit method*, Graduate Studies in Mathematics, vol. 64, American Mathematical Society, Providence, RI, 2004. MR 2069175

[Ser77] Jean-Pierre Serre, *Linear representations of finite groups*, Springer-Verlag, New York-Heidelberg, 1977, Translated from the second French edition by Leonard L. Scott, Graduate Texts in Mathematics, Vol. 42. MR 0450380

[Wil19] Dana P. Williams, *A tool kit for groupoid C^*-algebras*, Mathematical Surveys and Monographs, vol. 241, American Mathematical Society, Providence, RI, 2019. MR 3969970