Continuous representations

Bill Casselman
University of British Columbia
cass@math.ubc.ca

This essay is intended to be an introduction to the theory of continuous representations of locally compact groups. It contains somewhat dry material, mostly useful in motivating eventually a certain crucial but at first sight somewhat technical transition from representations of groups to representations of Lie algebras. Parts of it will also be used in the theory of automorphic forms. I have made some effort to reduce everything to well known facts in measure theory and topology.

The standard references for the material here is [Borel:1972] and [Weil:1965]. I have also used [Stein:1970].

Contents

This topic necessarily involves rather general **t**opological **v**ector **s**paces (as opposed to, say, only Hilbert spaces). In this paper, a TVS will always be assumed to be locally convex and Hausdorff. But as a rule, all topological vector spaces occurring in representation theory are also quasi-complete. Quasi-completeness is at first sight a rather technical condition, but in fact very practical. A quasi-complete TVS V is one for which integrals of V-valued functions are well defined, and in which derivatives can be characterized in a particularly useful way. Nearly all TVS encountered in the real world are quasi-complete, and it is rare that one has to think much about it. For example, Fréchet spaces and LF spaces are quasi-complete, and so are their duals. The standard references on quasi-complete spaces are [Treves:1967], §VI.5 of [Bourbaki:Integration], and §III.8 of [Bourbaki:TVS].

Note: I decline 'TVS' as I do 'sheep' or 'fish': 'one TVS', 'two TVS', etc. I wish to thank Murat Güngör for pointing out to me some gaps in an earlier exposition on smooth representations.

Throughout, let G be a locally compact topological group with a countable basis of neighbourhoods of 1, and which is a countable union of compact subsets. I assume it to be unimodular. Fix on it an invariant Haar measure dx. As a consequence of these assumptions, $\mathrm{L}^2(G)$ will be separable (Theorem 28.2 of [Hewitt-Ross:1969]).

1. Continuous representations

A continuous representation of G on a finite-dimensional space V is simply a continuous homomorphism from G to $\mathrm{GL}(V)$. If we are given a choice of coordinates, this means that matrix entries are continuous functions. But infinite-dimensional representations require more care. Suppose V to be a quasi-complete Hausdorff topological vector space. Following what we have just seen, we might say a representation of G is continuous if matrix coefficients $\langle \widetilde{v}, \pi(g)v\rangle$ are continuous for arbitrary v in V, $\widehat{v}$ in the continuous linear dual of V. In fact, for unitary representations on a Hilbert space, this is an adequate definition. But for arbitrary vector spaces V, complications arising because of a large number of semi-norms have to be taken into account.

A **continuous representation** of G on V is a homomorphism π from G to the group of linear transformations of V such that

$$G \times V \to V\colon\ (g, v) \longmapsto \pi(g)v$$

is continuous.

By considering an explicit matrix representation, one can see easily that this definition agrees with the naive one if V has finite dimension.

The defining condition means that whenever we are given $g_\circ$ in G, $v_\circ$ in V, and a semi-norm ρ of V then we can find a neighbourhood X of 1 in G and a semi-norm σ such that

$$\|\pi(xg_\circ)(v_\circ + u) - \pi(g_\circ)v_\circ\|_\rho < \varepsilon$$

whenever x lies in X and $\|u\|_\sigma < \delta$.

It is often annoying to check the definition directly, but verification can be reduced to two simpler steps.

1.1. Proposition. *The representation (π, V) is continuous if and only if these two conditions are satisfied:*

(a) for a fixed v in V the map $g \mapsto \pi(g)v$ is continuous;
(b) if X is a compact subset of G, v in V, and ρ a semi-norm of V, there exists a semi-norm σ such that

$$\|\pi(x)v\|_\rho \le \|v\|_\sigma \quad (x \in X)\ .$$

The first condition, in our circumstances, means

(a′) *Suppose $v_\circ$ in V. For every continuous semi-norm ρ and $\varepsilon > 0$ there exists a neighbourhood X of 1 in G such that $\|\pi(x)v_\circ - v_\circ\|_\rho < \varepsilon$ whenever x lies in X.*

The second condition (b) here is usually the easier to verify, since it says that the family of semi-norms on V is in some sense invariant under G, and this is often transparently true. In fact, under a mild restriction on V (that it be *barreled*) (b) follows from (a). This is to be found as Proposition 1 of §VIII.1 in Bourbaki's **Integration**.

Proof of the Proposition. The necessity of (a′) and (b) is immediate from the definition of continuity. As for sufficiency, suppose $g_\circ$, $v_\circ$, and ρ given. Then

$$\begin{aligned}\|\pi(xg_\circ)(v_\circ + u) - \pi(g_\circ)v_\circ\|_\rho &= \|\pi(xg_\circ)v_\circ - \pi(g_\circ)v_\circ + \pi(x)\pi(g_\circ)u\|_\rho \\ &\le \|\pi(x)v_1 - v_1\|_\rho + \|\pi(x)\pi(g_\circ)u\|_\rho \quad (v_1 = \pi(g_\circ)v_\circ)\,.\end{aligned}$$

According to (a′) we can find a neighbourhood X of 1 in G such that

$$\|\pi(x)v_1 - v_1\|_\rho < \varepsilon/2\,,$$

and then by (b) we can find σ such that

$$\|\pi(x)\pi(g_\circ)u\|_\rho \le \|u\|_\sigma < \varepsilon/2$$

for x in X, if $\|u\|_\sigma < \varepsilon/2$. ▯

One immediate consequence of the proposition:

1.2. Corollary. *If V has finite dimension, the representation (π, V) is continuous if and only if the map $\pi: G \to \mathrm{GL}(V)$ is continuous.*

That is to say, if and only if matrix entries are continuous functions on G.

The space $C(G, V)$ is that of all continuous functions on G with values in V. Its topology is defined by semi-norms

$$\|F\|_{\Omega,\rho} = \sup_{g\in\Omega} \|F(g)\|_\rho$$

where Ω is a compact subset of G and ρ a continuous semi-norm on V.

Given a continuous representation (π, V) there exists a canonical embedding of V into $C(G, V)$:

$$v \longmapsto [g \longmapsto \pi(g)v]\,.$$

The image of V is in fact the closed subspace of all functions F such that $F(gx) = \pi(g)F(x)$ for all g, x in G.

1.3. Corollary. *The representation (π, V) is a continuous representation of G if and only if this embedding of V into $C(G, V)$ is continuous.*

Proof. Condition (a) means that the image of V lies in $C(G, V)$. Condition (b) means that the map from V to $C(G, V)$ is continuous. ▯

If (π, V) is a continuous representation of G, then *a priori* we have two topologies on V, the original one and that induced from $C(G, V)$. Condition (b) reassures us that these are the same.

The simplest general class of continuous representations is that of representations of G on various spaces of functions on itself and on quotient spaces $H\backslash G$, as well as on certain spaces of induced representations. The group G acts on these by means of the **right regular**, and in some cases **left regular**, representations:

$$R_g F(x) := F(xg), \quad L_g F(x) := F(g^{-1}x)\,.$$

On G itself these commute, hence define an action of $G \times G$.

1.4. Example. I'll not give an exhaustive list of examples, but offer two types as models. They will all be spaces of functions on quotients $H\backslash G$ in which H is a closed subgroup of G.

I recall a bit of topology. Suppose X to be any locally compact Hausdorff space which is the union of a countable number of compact subsets. The TVS $C(X)$ is that of all continuous functions on X. For each compact subset Ω in X define the semi-norm

$$\|f\|_\Omega = \sup_{x\in\Omega} |f(x)|\,.$$

These make X into a Fréchet space.

The TVS $C_c(X)$ is that of all continuous functions on X whose support is a relatively compact open subset of X. If Y is a relatively compact open subset, let $C_c(\overline{Y})$ be the subspace of all f in $C(\overline{Y})$ that vanish on the boundary of Y. Such a function may be extended uniquely to a function in $C(X)$ vanishing outside Y. Define the semi-norm

$$\|f\|_Y = \sup_{x \in Y} |f(x)| \, .$$

This makes $C_c(X)$ into an LF space.

Both spaces are quasi-complete, as are their duals.

1.5. Proposition. *Suppose H a closed subgroup of G. The right regular representations of G on $C(H\backslash G)$ and $C_c(H\backslash G)$ are continuous.*

Proofs should be evident.

1.6. Example. Suppose G to be unimodular, H a closed subgroup. Let δ_H be the modulus character of H, and let Ω_c (for real $0 \le c \le 1$) be the space of all continuous functions $F\colon G \to \mathbb{C}$ such that $F(hg) = \delta_H^c(h) f(g)$ for all h in H, g in G. and which are of compact support modulo H. There exists on Ω_c a positive, continuous, G-invariant linear functional, which I'll express as integration. This may be embedded in a space of integrable elements Ω_1. Let $\Omega_{1/2}$ be the space of functions on G such that $f(hg) = \delta_H^{1/2}(h) f(g)$. This implies that f^2 lies in Ω_1.

1.7. Proposition. *The right-regular representations of G on $\Omega_1(H\backslash G)$ and $\Omega_{1/2}(H\backslash G)$ are continuous.*

If (π, V) is a continuous representation of G then its dual $(\widehat{\pi}, \widehat{V})$ on the continuous linear dual of V is that defined by the condition

$$\langle \pi(g)v, \widehat{\pi}(g)\widehat{v} \rangle = \langle \widehat{v}, v \rangle \; .$$

In other words,

$$\langle v, \widehat{\pi}(g)\widehat{v} \rangle = \langle \pi(g)^{-1} v, \widehat{v} \rangle \; .$$

The continuous linear dual of a TVS may be assigned the weak topology, with semi-norms $\|\widehat{v}\|_v = |\langle v, \widehat{v} \rangle|$ for v in V. It is straightforward to prove:

1.8. Proposition. *If (π, V) is a continuous representation of G, the dual representation $\widehat{\pi}$ is continuous in the weak topology on $\widehat{V}$.*

The **matrix coefficient** corresponding to the pair $\widehat{v}$ and v is the continuous function

$$\mu_{\hat{v},v}\colon \; g \longmapsto \langle \pi(g^{-1})v, \widehat{v} \rangle$$

on G. If π is finite-dimensional and (e_i) is a basis of V with dual basis $\widehat{e}_i$, then $\langle e_i, \widehat{e}_j \rangle$ is in fact a matrix entry.

2. Representation of measures

Suppose (π, V) to be a continuous representation of G on a quasi-complete topological vector space.

THE OPERATORS. For us, the basic property of quasi-complete topological vector spaces V is that if μ is a bounded measure of compact support Ω on a locally compact space X and f is a continuous function on X with values in V then the integral

$$I_f = \int_X f(x)\, d\mu(x)$$

has a well defined value in V, which is characterized uniquely by the property

$$\langle \widehat{v}, I_f \rangle = \int_X \langle \widehat{v}, f(x) \rangle \, d\mu(x)$$

for all $\widehat{v}$ in $\widehat{V}$. If ρ is a (continuous) semi-norm of V then

$$\big\| I_f \big\|_\rho \le \sup_\Omega \big\| f(x) \big\|_\rho \cdot \int_\Omega d|\mu| \, .$$

For this material, I refer to Chapter VI of [Bourbaki:1959-65].

The space $M_c(G)$ can be identified with the space of continuous linear functionals on the space $C(G, \mathbb{C})$ of all continuous functions on G. Given the representation π, every μ in $M_c(G)$ determines an operator on V:

$$\pi(\mu)v = \int_G \pi(g)v \; d\mu(g) \, .$$

Suppose Ω to be the support of μ. Condition (b) of Proposition 1.1 implies that for every semi-norm ρ of V there exists a semi-norm σ such that

$$\|\pi(g)v\|_\rho \le \|v\|_\sigma$$

for all g in Ω. Then

(2.1)
$$\begin{aligned} \big\|\pi(\mu)v\big\|_\rho &\le \int_\Omega \big\|\pi(g)v\big\|_\rho \, |d\mu| \\ &\le \left(\int_\Omega |d\mu| \right) \|v\|_\sigma \, . \end{aligned}$$

Hence:

2.2. Proposition. *For every μ in $M_c(G)$ the operator $\pi(\mu)$ is continuous.*

There are two actions of G on $M_c(G)$, left L and right R. We have

$$\pi(L_g\mu) = \int_G \pi(x)\, dL_g\mu = \int_G \pi(gx)\, d\mu(x) = \pi(g)\pi(\mu) \, .$$

For a given v in V the map $\mu \mapsto \pi(\mu)v$ from $M_c(G)$ to V is equivariant with respect to this left regular representation. This allows us to identify V with $\mathrm{Hom}_G(M_c(G), V)$ as spaces. But the right regular representation of G assigns it a compatible G-structure. Thus:

2.3. Corollary. *This map identifies V with the G-space $\mathrm{Hom}_G(M_c(G), V)$.*

THE RING OF MEASURES. The space $M_c(G)$ possesses a canonical ring structure. If μ_1 and μ_2 are two measures in $M_c(G)$, the tensor product $\mu_1 \otimes \mu_2$ is the unique measure on $G \times G$ such that

$$\langle \mu_1 \otimes \mu_2, X_1 \times X_2 \rangle = \langle \mu_1, X_1 \rangle \langle \mu_2, X_2 \rangle \,.$$

This measure may be identified with a continuous linera function on $C(G \times G)$. For f in $C(G \times G)$ and x in G, the integral

$$\int_G f(x,y)\, d\mu_1(x)$$

is a continuous function on G, and according to Fubini's Theorem the product integral is

$$\int_G \left(\int_G f(x,y)\, d\mu_1(x) \right) d\mu_2(y) \,.$$

The product of two measures in $M_c(G)$ product is their convolution, the measure defined by the formula

$$\langle \mu_1 * \mu_2, f \rangle = \int_{G \times G} f(xy)\, d(\mu_1(x) \otimes \mu_2(y)) \,.$$

The multiplicative identity in $M_c(G)$ is the Dirac δ_1 taking f to $f(1)$.

Given the Haar measure dx, the space $C_c(G)$ of continuous functions with compact support may be embedded in $M_c(G)$: $f \mapsto f\, dx$.

If $\mu_1(x) = F_1(x)\, dx$ the convolution integral is

$$\begin{aligned} \int_G \left(\int_G f(xy) F_1(x)\, dx \right) d\mu_2(y) &= \int_G \left(\int_G f(z) F(zy^{-1})\, dz \right) d\mu_2 \\ &= \int_G f(z) \left(\int_G F(zy^{-1})\, d\mu_2(y) \right) dz \,. \end{aligned}$$

One consequence:

2.4. Proposition. *If $\mu_1(x) = F(x)\, dx$ then $\mu_1 * \mu_2$ is the continuous function*

$$\int_G F(zy^{-1})\, d\mu_2(y) \,.$$

For any two F_i in $C_c(G)$ define

$$[F_1 * F_2](y) = \int_G F_1(x) F_2(x^{-1}y)\, dx$$

Another consequence:

2.5. Proposition. *If $\mu_i = F_i\, dx$ then*

$$\mu_1 * \mu_2 = (F_1 * F_2)\, dx \,.$$

2.6. Proposition. *The operators $\pi(\mu)$ make V a module over the ring $M_c(G)$:*

$$\pi(\mu_1)\pi(\mu_2) = \pi(\mu_1 * \mu_2) \ .$$

DIRAC MEASURES. I define a **Dirac function** on G to be a function f in $C_c^\infty(G, \mathbb{R})$ of the form

$$f = h * h^\vee$$

with $h \ge 0$, such that

$$\int f(g)\, dg = 1 \, .$$

If f is a Dirac function with support in the compact neighbourhood Ω of 1 then

$$\pi(f)v - v = \int_\Omega f(g)(\pi(g)v - v)\, dg \, .$$

and

$$\|\pi(f)v - v\|_\rho \le \mathrm{meas}(\Omega) \cdot \sup_{g \in \Omega} \|\pi(g)v - v\|_\rho \, .$$

Hence:

2.7. Proposition. *Suppose (f_n) to be a sequence of Dirac functions whose support has limit $\{1\}$. Then $\lim_n \pi(f_n)v = v$ for any v in V.*

3. Representations of a compact group I. Finite-dimensional

Representations of compact groups are models for all of representation theory and are also required as a necessary preliminary to much of it. In this section, let K be an arbitrary compact group with a countable basis of neighbourhoods of 1. I do not require, at least at first, that it be a Lie group. For example, the group $\mathrm{GL}_n(\mathbb{Z}_p)$, which is compact according to König's Lemma, is allowable.

I make no claim to originality. The exposition pretty much follows a straightforward line, but because it is rather long I have divided it into two major pieces, mostly separating finite dimensions from infinite. The theory of continuous finite dimensional (**cfd**) representations of compact groups differs from the theory of representations of finite groups primarily only in so far as integration replaces finite sums, so this first part is especially simple. I have included some very elementary results about linear algebra in an appendix.

I'll begin by discussing only finite-dimensional representations, but I should say right here that every irreducible continuous representation of a compact group is finite-dimensional. This will be proved much later (as Corollary 4.13).

3.1. Proposition. *The group K is unimodular.*

This means that every left-invariant Haar measure is also right-invariant.

Proof. A left-invariant Haar measure $d_\ell x$ is unique up to constants. For k in K, $d_\ell xk$ is also left-invariant, hence $d_\ell xk = \delta(k)\, d_\ell x$ for some positive constant $\delta(k)$. The map taking k to $\delta(k)$ is a continuous homomorphism from K to the multiplicative group of positive real numbers, hence trivial. Therefore $d_\ell xk = d_\ell x$, so $d_\ell x$ is also right-invariant. ∎

- *Assign K, once and for all, an invariant measure dk of total measure 1.*

This allows us to identify continuous functions on K with measures: $f \mapsto f\,dk$.

SEMI-SIMPLICITY. Define $\mathbf{1}$ to be the Haar measure on K, the measure associated to the constant function equal to 1 on all of K. If (π, V) is a continuous representation of K then

$$\pi(\mathbf{1})v = \int_K \pi(k)v\,dk\,.$$

The vector $\pi(\mathbf{1})v$ is fixed by all of K. If v is already fixed by K then $\pi(\mathbf{1})v = v$, so $\pi(\mathbf{1})$ is idempotent.

3.2. Proposition. *If (π, V) is a continuous representation of K then v is fixed by all of K if and only if $\pi(\mathbf{1})v = v$. The kernel of $\pi(\mathbf{1})$ is a closed K-stable summand of V and*

$$V = V^K \oplus \operatorname{Ker} \pi(\mathbf{1})\,.$$

In other words, $\pi(\mathbf{1})$ is the unique K-equivariant projection onto the subspace V^K of K-fixed vectors, which is a closed subspace of V.

Proof. We have

$$v = \int_K \pi(k)v\,dk + \int_K \big(I - \pi(k)\big)v\,dk\,.$$

The first term is $\pi(\mathbf{1})v$ and is fixed by K. The second is in the kernel of $\pi(\mathbf{1})$. ▯

An application of Hahn-Banach will show that $\operatorname{Ker} \pi(\mathbf{1})$ is the closure of the span of the $\pi(k)v - v$.

3.3. Corollary. *If*

$$0 \longrightarrow U \longrightarrow V \longrightarrow W \longrightarrow 0$$

is an exact sequence of continuous representations of K then

$$0 \longrightarrow U^K \longrightarrow V^K \longrightarrow W^K \longrightarrow 0$$

is also exact.

One major point of this section is to find a generalization of the projection $\pi(\mathbf{1})$ associated to representations of K other than the trivial one.

3.4. Proposition. *Any short exact sequence*

$$0 \longrightarrow U \longrightarrow V \longrightarrow W \longrightarrow 0$$

of continuous K-representations, where W has finite dimension, splits continuouslyand equivariantly.

Thus $V \cong W \oplus U$ as a K-module.

Proof. Choose any linear splitting $f \colon W \to V$, which is necessarily continuous. Then

$$\overline{f} \colon\ w \longmapsto \int_K \pi(k)\, f\big(\pi(k^{-1})w\big)\, dk$$

is a K-equivariant splitting. ▯

3.5. Corollary. *Every cfd representation of K is a direct sum of irreducible representations.*

This follows by induction from Proposition 3.4.

3.6. Corollary. *Suppose (π, V) to be a cfd representation of K. It is irreducible if and only if $\mathrm{End}_K(V) = \mathbb{C}$.*

That is to say that it is irreducible if and only if it is indecomposable.

Proof. If π is reducible, $\mathrm{End}_K(V)$ contains at least two linearly independent projection operators.

Now say π is irreducible, and suppose F is a non-trivial K-equivariant endomorphism. If c is an eigenvalue, then $\mathrm{Ker}(I - cI)$ is K-invariant, hence must be all of V. ▯

PRODUCTS. If (π_1, V_1) and (π_2, V_2) are cfd representations of K_1 and K_2, the tensor product representation of $K_1 \times K_2$ is defined by the formula

$$[\pi_1 \otimes \pi_2](k_1 \times k_2)\colon\ v_1 \otimes v_2 \longmapsto \pi_1(k_1)v_1 \otimes \pi_2(k_2)v_2\ .$$

3.7. Proposition. *If K_1, K_2 are compact groups and (π_i, V_i) are irreducible cfd representations of K_i, then $\pi_1 \otimes \pi_2$ is irreducible.*

Proof. It must be shown that

$$\mathrm{End}_{K_1 \times K_2}(V_1 \otimes V_2) = \mathbb{C}\,.$$

Well,

$$\mathrm{End}(V_1 \otimes V_2) = \mathrm{End}(V_1) \otimes \mathrm{End}(V_2)$$

and the subring of $K_1 \times K_2$ invariants is

$$\mathrm{End}_{K_1}(V_1) \otimes \mathrm{End}_{K_2}(V_2)\,.$$ ▯

3.8. Proposition. *Every irreducible cfd representation of $K_1 \times K_2$ is the tensor product of irreducible representations of K_1 and K_2.*

Proof. Say (π, V) to be an irreducible representation of $K_1 \times K_2$. Suppose $U \neq 0$ to be any irreducible K_1-stable subspace. This means that we are given a non-trivial map in $\mathrm{Hom}_{K_1}(U, V)$. The canonical map

$$U \otimes \mathrm{Hom}_{K_1}(U, V) \longrightarrow V$$

is therefore non-trivial. Hence there exists some irreducible representation W of K_2 on a space V_2 for which

$$U \otimes W \longrightarrow V$$

is non-trivial. Here $U \otimes W$ is an irreducible representation of $K_1 \times K_2$, according to Proposition 3.7. Because V is irreducible, this map is an isomorphism. ▯

SCHUR'S LEMMA. The following is an immediate extension of Corollary 3.6:

3.9. Proposition. (Schur's Lemma) *Suppose π_1, π_2 to be two irreducible cfd representations of K. Then*

$$\mathrm{Hom}_K(V_1, V_2) = \begin{cases} 0 & \text{if } \pi_1 \text{ is not isomorphic to } \pi_2. \\ \mathbb{C} & \text{if } \pi_1 = \pi_2 \end{cases}$$

If the (π_i, V_i) are irreducible cfd representations of K, then $V_1 \otimes \widehat{V}_2$ is isomorphic to $\mathrm{Hom}(V_2, V_1)$. If π_1 and π_2 are not isomorphic, then the subspace of K-invariants is trivial. If they are equal, the invariant endomorphisms is made up of scalar maps. Hence:

3.10. Corollary. *Suppose π_1, π_2 to be two irreducible cfd representations of K. Then the subspace of K-invariants in $V_1 \otimes V_2$ is trivial unless π_2 is isomorphic to the dual of π_1, and then it consists of scalar multiples of the canonical pairing.*

Explicitly, if $\{e_i\}$ is a basis of V and $\{\widehat{e}_i\}$ the dual basis, then

$$\sum e_i \otimes \widehat{e}_i$$

spans the K-invariants in $V \otimes \widehat{V}$.

MATRIX COEFFICIENTS. The product $K \times K$ acts on $C(K)$, the space of continuous functions on K, by

$$[\rho_{k_1 \times k_2} f](k) = [L_{k_1} R_{k_2} f](k) = f(k_1^{-1} k k_2)\,.$$

If (π, V) is a cfd representation, the **matrix coefficient** associated to a pair $(v, \widehat{v})$ is the continuous function

$$\mu_{v \otimes \widehat{v}}(k) = \langle \pi(k^{-1})v, \widehat{v} \rangle\,.$$

Matrix coefficients define a $K \times K$-equivariant map

$$V \otimes \widehat{V} \longrightarrow C(K),\ v \otimes \widehat{v} \longmapsto \mu_{v \otimes \widehat{v}}\,.$$

Swapping the roles of V and $\widehat{V}$, we get also a second equivariant map

$$\widehat{V} \otimes V \longrightarrow C(K),\ \widehat{v} \otimes v \longmapsto \mu_{\widehat{v} \otimes v}\,.$$

We can then define an invariant pairing of two of these spaces by integrating the product. Assume π to be irreducible, both $\pi_i = \pi$. How does this pairing of $V \otimes \widehat{V}$ with $\widehat{V} \otimes V$ compare with the natural one? The space $V \otimes \widehat{V}$ is an irreducible representation of $K \times K$. It is isomorphic to $\mathrm{End}(V)$, and hence according to Proposition 3.9 the space of invariants has dimension one. Therefore one must be a scalar multiple of the other. In fact:

3.11. Proposition. (Schur orthogonality, matrix coefficients) *Suppose π_1, π_2 to be two irreducible representations of K. For v_i, $\widehat{v}_i$ in V_i, $\widehat{V}_i$*

$$\int_K \langle \pi_1(k^{-1})v_1, \widehat{v}_1 \rangle \langle \pi_2(k) v_2, \widehat{v}_2 \rangle \, dk = \begin{cases} 0 & \text{unless } \pi_1 \text{ is isomorphic to } \pi_2 \\ \dfrac{1}{\dim \pi} \langle v_1, \widehat{v}_2 \rangle \langle v_2, \widehat{v}_1 \rangle & \text{if } \pi = \pi_1 = \pi_2 \end{cases}$$

Proof. It remains only to determine the constant c_π such that

$$\int_K \langle \pi(k)v_1, \widehat{v}_1 \rangle \langle \pi(k^{-1}) v_2, \widehat{v}_2 \rangle \, dk = c_\pi \, \langle v_2, \widehat{v}_1 \rangle \langle v_1, \widehat{v}_2 \rangle\,.$$

Choose a basis (e_i) and its dual $(\widehat{e}_i)$. The equation above implies that

$$\int_K \sum_j \langle \pi(k)e_j, \widehat{e}_i\rangle\langle \pi(k^{-1})e_\ell, \widehat{e}_j\rangle\, dk = c_\pi \cdot \sum_j \langle e_\ell, \widehat{e}_i\rangle\langle e_j, \widehat{e}_j\rangle$$
$$= c_\pi \dim(\pi)\cdot\langle e_\ell, \widehat{e}_i\rangle\,.$$

But the integrand is the (i,ℓ) entry of $\pi(k)\pi(k^{-1}) = I$, so $c_\pi \cdot \dim(\pi) = 1$.

There are other versions of this to come.

HERMITIAN FORMS. A **Hermitian pairing** between two vector spaces is a map

$$v_1 \otimes v_2 \longmapsto v_1 \bullet v_2$$

from $V_1 \times V_2$ to $\mathbb{C}$, linear in the first factor, conjugate linear in the second, and symmetric in the sense that

$$v_1 \bullet v_2 = \overline{v_2 \bullet v_1}\,.$$

A **Hermitian form** on V is a Hermitian pairing of V with itself. If H is a Hermitian form on V and c is in $\mathbb{C}$, then cH is Hermitian if and only if c is real. The set of all positive definite Hermitian forms on V is a convex open cone in the space of all Hermitian forms on V.

If V is a complex vector space, the **conjugate** vector space is the same space, but with the conjugate scalar multiplication:

$$c\bullet v = \overline{c}\cdot v\,.$$

In these terms, a Hermitian form is a linear map from V to the conjugate of its linear dual.

3.12. Proposition. *If (π, V) is an irreducible cfd representation, then the space of K-invariant Hermitian forms on V has dimension one over $\mathbb{R}$. The subset of positive definite ones is a real ray.*

The last assertion means that there exists at least one positive definite Hermitian form on V that is K-invariant, and all others are positive multiples of it. Consequently, every cfd representation is **unitary**.

Proof. The only non-trivial point to verify is that there always exists a positive definite invariant Hermitian form on V. But if one starts with an arbitrary positive definite form on V its K-average will still be positive definite, because of convexity.

3.13. Proposition. *If (π_1, V_1) and (π_2, V_2) are two irreducible cfd representations, the space of K-invariant Hermitian pairings of V_1 with V_2 is trivial if π_1 is not isomorphic to π_2.*

Proof. Any Hermitian pairing is a conjugate-linear map from V_1 to $\widehat{V}_2$. Because of irreducibility, it must be an isomorphism. But the previous result tells us that there exist conjugate-linear K-isomorphism of V_2 with $\widehat{V}_2$, hence a linear K-isomorphism of V_1 with V_2.

Suppose (π, V) to be an irreducible cfd representation, $u\bullet v$ an invariant positive definite Hermitian form on V. Since it is non-degenerate, for each $\widehat{v}$ in $\widehat{V}$ there exists a vector $\varphi(\widehat{v})$ in V such that

$$v\bullet\varphi(\widehat{v}) = \langle v, \widehat{v}\rangle$$

for all v in V. The map $\widehat{v} \mapsto \varphi(\widehat{v})$ is conjugate-linear and K-equivariant. We can assign a positive definite Hermitian form on $\widehat{V}$ by the formula

$$\widehat{u}\,\widehat{\bullet}\,\widehat{v} = \varphi(\widehat{v})\bullet\varphi(\widehat{u})\,,$$

which is conjugate-linear in the second factor. How does this depend on the choice of Hermitian inner product? If we are given two Hermitian inner products $\bullet_1$ and $\bullet_2 = c \cdot \bullet_1$, then

(3.14)
$$\widehat{\bullet}_2 = c^{-1} \cdot \widehat{\bullet}_1 \,.$$

Given a Hermitian form on V, one can then define one on $V \otimes \widehat{V}$ according to the formula

$$(v_1 \otimes \widehat{v}_1) \bullet (v_2 \otimes \widehat{v}_2) = (v_1 \bullet v_2) \cdot (\widehat{v}_1 \,\widehat{\bullet}\, \widehat{v}_2) \,.$$

Because of (3.14) does not change if φ is scaled.

The representation of $K \times K$ on $C(K)$ extends to a continuous representation of $K \times K$ on the Hilbert space $\mathrm{L}^2(K)$ with the invariant Hilbert norm

$$f_1 \bullet f_2 = \int_K f_1(k) \overline{f_2(k)} \, dk \,.$$

The canonical positive definite Hermitian form on $V \otimes \widehat{V}$ and the Hilbert norm restricted to the matrix coefficient image must be scalar multiples of each other. Explicitly:

3.15. Corollary. (Unitary Schur orthogonality) *Suppose π_1, π_2 to be two irreducible representations of K. For v_i, $\widehat{v}_i$ in V_i, $\widehat{V}_i$*

$$\int_K \langle \pi(k^{-1}) v_1, \widehat{v}_1 \rangle \overline{\langle \pi(k^{-1}) v_2, \widehat{v}_2 \rangle} \, dk = \begin{cases} 0 & \text{unless } \pi_1 \text{ is isomorphic to } \pi_2 \\ \dfrac{1}{\dim \pi} (v_1 \otimes \widehat{v}_1) \bullet (v_2 \otimes \widehat{v}_2) & \text{if } \pi = \pi_1 = \pi_2 \end{cases}$$

A TRACE FORM OF SCHUR'S FORMULA. As explained in the appendix, there is a canonical equivariant map from $V \otimes \widehat{V}$ to the ring of endomorphisms of V_π: $v \otimes \widehat{v}$ corresponds to the endomorphism

$$\varepsilon_{v \otimes \widehat{v}} \colon u \longmapsto \langle u, \widehat{v} \rangle v \,.$$

The image of this map is the ring of all linear transformations of finite rank.

What map from $\mathrm{End}(V)$ to $C(K)$ corresponds to the matrix coefficient map μ? Define $\tau = \tau_\pi$ from $\mathrm{End}(V_\pi)$ to $C(K)$:

$$\tau_\pi(F) \colon k \longmapsto \mathrm{trace}(\pi(k^{-1}) \, F) \,.$$

Then

3.16. Lemma. *The diagram*

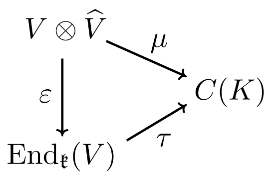

is commutative.

Proof. The trace of $\varepsilon_{v \otimes \widehat{v}}$ is $\langle v, \widehat{v} \rangle$. Hence if $F = \varepsilon_{v \otimes \widehat{v}}$ the trace of $\pi(k^{-1}) F$ is $\langle \pi(k^{-1}) v, \widehat{v} \rangle$. ∎

This observation leads to yet another form of Schur's relations. If $T = (t_{i,j})$ is in $\mathrm{End}(V)$, its **adjoint** T^* is defined by the equation

$$T^*u\bullet v = u\bullet Tv\,.$$

In matrix terms, it is the transposed conjugate. Then

$$T\bullet T = \mathrm{trace}(T\cdot T^*) = \sum |t_{i,j}|^2$$

is a $G\times G$-invariant positive definite Hermitian norm on $\mathrm{End}(V)$.

Here is yet another version of Schur's formula:

3.17. Corollary. (Schur's formula, trace form) *For E in $\mathrm{End}(V_\pi)$, F in $\mathrm{End}(V_\rho)$*

$$\int_G \mathrm{trace}(\pi(k^{-1})\cdot E)\cdot\overline{\mathrm{trace}(\rho(k^{-1})\cdot F)}\,dk = \begin{cases} 0 & \text{if } \pi \text{ is not isomorphic to } \rho \\ \dfrac{1}{\dim\pi}\,E\bullet F & \text{if } \pi=\rho\,. \end{cases}$$

Since $\pi(k)$ is unitary, the trace of $\pi(k^{-1})$ is the conjugate of that of $\pi(k)$. This is hence consistent with Proposition 3.11.

3.18. Corollary. *For two representations π, ρ:*

$$\int_K \mathrm{trace}(\pi(k^{-1}))\cdot\overline{\mathrm{trace}(\rho(k^{-1}))}\,dk = \begin{cases} 0 & \text{if } \rho \text{ is not isomorphic to } \pi \\ 1 & \text{if } \pi=\rho. \end{cases}$$

THE OPERATOR FORM. Other versions of Schur's formulas lead to the Hilbert space version of Plancherel's formula, the following to a formula for the inverse Fourier transform. I recall the map

(3.19)
$$C(K)\longrightarrow \mathrm{End}(V),\ f\longmapsto \pi(f)\,.$$

3.20. Theorem. *Suppose (π,V), (ρ,U) to be irreducible cfd representations of K. The composition*

$$\mathrm{End}(V_\pi)\xrightarrow{\ \tau\ } C(K)\xrightarrow{\ \rho\ }\mathrm{End}(V_\rho)$$

amounts to scalar multiplication by

$$\begin{cases} 0 & \text{if } \pi \text{ is not isomorphic to } \rho; \\ 1/\dim(\pi) & \text{if } \pi=\rho. \end{cases}$$

Proof. Both left and right are in $\mathrm{End}(V)$, and the map from right to left commutes with the action of $G\times G$, hence the two sides differ by a scalar. This can be evaluated by setting $f=I$.
▯

Explicitly, this says

$$\int_K \mathrm{trace}\,(\pi(k^{-1})f)\,\pi(k)\,dk = \frac{1}{\dim\pi}\,f$$

for f in $\mathrm{End}(V_\pi)$. In a way, this formula is a precursor of Corollary 3.17.

There is an important special case of this. For each irreducible cfd representation (π, V) define the function

$$\xi_\pi(k) = \dim(\pi)\cdot\mathrm{trace}(\pi(k^{-1}))\,.$$

in $C(K)$.

3.21. Theorem. *Suppose (π, V) to be an irreducible cfd representation of K.*

(a) *if π is irreducible then $\pi(\xi_\pi) = I$;*
(b) *if π, ρ are two non-isomorphic irreducible cfd representations of K then $\pi(\xi_\rho) = 0$.*

Suppose V to be a cfd representation. It V will be a direct sum of irreducible cvds with multiplicity. If (σ, U) is an irreducible cfds, then $\pi(\xi_\sigma)$ will amount to projection to the sum of copies of U, the **isotypic** σ-component. It is also the canonical image of the embedding of $U \otimes \mathrm{Hom}(U, V)$ into V. Every equivariant map from σ factors through it. So ξ_σ is the replacement for the projection $\pi(\mathbf{1})$.

3.22. Theorem. *Suppose (π, V) to be an irreducible cfd representation of K. Then*

(a) *the function ξ_π lies in the centre of $C(K)$;*
(b) *it is idempotent with respect to convolution.*

Proof. The first because trace is conjugate-invariant. The second from Corollary 3.17 with $F = I$. ▯

Finally, the generalization of Proposition 3.2:

3.23. Corollary. *Suppose (π, V) to be any continuous representation of K, σ an irreducible cfd, and $\xi = \xi_\sigma$. Let V^ξ be the space of vectors fixed by ξ and V_ξ its kernel. Then*

$$V = V^\xi \oplus V_\xi\,.$$

4. Representations of a compact group II. Infinite-dimensional

The main result of this section is the Peter-Weyl theorem—the orthogonal decomposition of $\mathrm{L}^2(K)$. Alongside of a few consequenes.

FROBENIUS RECIPROCITY. Let H be a closed subgroup of K. As we shall see, the Hilbert space $\mathrm{L}^2(K/H)$ is the orthogonal sum of irreducible representations of K of finite dimension, each with finite multiplicity. Which representations of K appear?

Define

$$\Omega\colon\ C(K/H) \longrightarrow \mathbb{C},\ f \longmapsto f(1)\,.$$

Then $\langle \Omega, L_h f\rangle = \langle \Omega, f\rangle$ for every h in H, which is to say that Ω is an H-invariant linear functional on $C(K/H)$.

4.1. Lemma. (Frobenius reciprocity) *Suppose (π, V) to be a cfd representation of K. Composition with Ω is an isomorphism of $\mathrm{Hom}_K(V, C(K/H))$ with $\mathrm{Hom}_H(V, \mathbb{C})$.*

Proof. Suppose

$$F\colon v \longmapsto F_v$$

to be a K-equivariant map from V to $C(K/H)$. If $F_v(1) = 0$ for all v in V, then

$$F_{\pi(g)v}(1) = L_g F_v(1) = F_v(g^{-1}) = 0$$

for all g in G, so $F_v = 0$.

Conversely, given an H-invariant map f from V to $\mathbb{C}$, each function $F_v(g) = f(\pi(g^{-1})v)$ is a K-equivariant map from V to $C(K/H)$.

I'll apply this to two cases:

(a) (π, V) is an irreducible cfd of K, $H = I$. In this case $C(K/H) = C(K)$ and the conclusion is that

$$\mathrm{Hom}_K(V, C(K) \cong \widehat{V}\,.$$

(b) (σ, U) is an irreducible cfd, $V = U \otimes \widehat{U}$, H is the diagonal in $K \times K$. Here the conclusion is that there is a $K \times K$-equivariant embedding from $U \otimes \widehat{U}$ to $C(K)$, which is unique up to scalar multplication.

Of course these are not exactly independent.

K-FINITE VECTORS. Suppose (π, V) to be any continuous representation of K. A **K-finite** vector in V is one contained in a finite-dimensional K-stable subspace of V.

4.2. Lemma. *Suppose f in $C(K)$. The following are equivalent:*

(a) f is K-finite with respect to the left regular representation;
(b) f is K-finite with respect to the right regular representation;
(c) f is K-finite with respect to both left and right regular representations.

In other words, in $C(K)$ the term 'K-finite' is unambiguous.

Proof. It suffices to show that (a) implies (c). Suppose given a K-stable subspace of $C(K)$ containing v. It suffices to assume that it is the image of some irreducible (σ, U). According to case (a) above, the embedding of U is of the form

$$u \longmapsto \langle \pi(g^{-1})u, \widehat{u}\rangle$$

for some $\widehat{u}$ in $\widehat{U}$. But then v is contained in the image of all of $U \otimes \widehat{U}$.

4.3. Lemma. *Every K-finite measure F on K is actually a continuous function.*

Proof. Suppose the measure F is contained in the finite-dimensional right-K-stable subspace U. Accordng to the previous Lemma, we may assume it to be some right- and left-stable cfd $U \otimes \widehat{U} = \mathrm{End}(U)$ with U irreducible. The image of $C(K)$ in $\mathrm{End}_{\mathbb{C}}(U)$ is closed, but by Proposition 2.7 the identity operator I is in its closure. That image therefore contains I, so for some f in $C(K)$ we have $R_f F = F$. But Proposition 2.4 tells us that if f lies in $C(K)$ then $R_f F$ is a continuous function.

THE PROJECTION OPERATORS. For any f in $C(K)$, let $f^* = \overline{f}^{\vee}$. The following is a basic and well known fact.

4.4. Lemma. *If $f = h*h^*$, then R_f and L_f, acting on $\mathrm{L}^2(K)$, are compact self-adjoint operators.*

They are in fact Hilbert-Schmidt operators. With non-negative eigenvalues, too, since (for example)

$$R_f\varphi\bullet\varphi = R_h\varphi\bullet R_h\varphi\,.$$

For the moment fix (σ, U) and let $\xi = \xi_\sigma$. Then $\xi = \xi * \xi$ and $\xi = \overline{\xi}^\vee$. According to the Lemma, R_ξ is a self-adjoint compact operator. The space $V = \mathrm{L}^2(K)$ therefore decomposes into the orthogonal sum of its kernel and a countable number of finite-dimensional eigenspaces. Since ξ is idempotent, so is R_f. There is exactly one such eigenspace, with eigenvalue 1. The corresponding eigenfunctions are certainly K-finite, and by Lemma 4.3 this space of eigenfunctions is the same as the embedded copy of $U \otimes \widehat{U}$. We have proved:

4.5. Proposition. *Suppose σ to be in irreducible cvd representation. Let $V = \mathrm{L}^2(K)$, $\xi = \xi_\sigma$. Then V^ξ is equal to the canonical copy of $V_\sigma \otimes \widehat{V}_\sigma$.*

The space $C(K)^\xi$ is thus isomorphic to $U \otimes \widehat{U}$. The space $C(K)$ is therefore the direct sum of $U \otimes \widehat{U}$ and the kernel of ξ. Hence. by duality (applied to $\widehat{\sigma}$):

4.6. Lemma. *The space $M(K)$ is isomorphic to the direct sum of $U \otimes \widehat{U}$ and the kernel of ξ.*

4.7. Proposition. *The natural map*

(4.8) $$U_\sigma \otimes \mathrm{Hom}_K(U_\sigma, V) \longrightarrow V^\xi\,.$$

is an isomorphism.

Proof. We know that $\mathrm{Hom}_K(M(K), V)$ is canonically isomorphic to V. The map takes v to the map $\mu \mapsto \pi(\mu)v$. But by Lemma 4.6

$$M(K) = M(K)^{\xi_\sigma} \oplus \mathrm{Ker}_{M(K)}(\xi_\sigma)\,.$$

Since $\xi_\sigma = 0$ on the second summand,

$$U_\sigma \otimes \mathrm{Hom}_K(U_\sigma, V) = \mathrm{Hom}_K(\widehat{U}_\sigma \otimes U_\sigma, \pi(\xi_\sigma V) = \pi(\xi_\sigma)V\,.$$

THE ORTHOGONAL DECOMPOSITION. Suppose (ρ, U) to be an irreducible cfd representation of K, $\xi = \xi_\rho$. Then L_ξ is the projection of $V = \mathrm{L}^2(K)$ onto the σ-component of V, which is isomorphic to $\mathrm{End}(U)$ as a module over $K \times K$. Fpr every f in $\mathrm{L}^2(K)$, let $f_\sigma = L_\xi f$. The following is a straightforward computation:

4.9. Lemma. *For any f in $C(K)$*

$$f_\xi = \dim(\sigma)\cdot\tau_\sigma(\sigma(f))\,,$$

Note that $\sigma(f)$ makes sense for f in $\mathrm{L}^2(K)$, since one can integrate the matrix entry by entry.

4.10. Theorem. (Peter-Weyl) *Every function in $\mathrm{L}^2(K)$ may be expressed as an orthogonal sum*

$$f = \sum\nolimits_\pi f_\sigma\,,$$

and

$$\|f\|^2 = \sum\nolimits_\sigma \dim(\sigma)\cdot\|\sigma(f)\|^2\,.$$

Proof. In several steps.

Step 1. It suffices to show that if φ in $\mathrm{L}^2(K)$ is orthogonal to all K-finite functions, it must vanish. So suppose φ to satisfy this condition. Then all $L_f\varphi$ for f in $C(K)$ also satisfy it, since the subspace of K-finite functions is stable under K, and if φ does not vanish then there exists f such that $L_f\varphi$ also doesn't vanish. But $R_f\varphi$ is in $C(K)$, so we may assume from now on that φ itself is orthogonal to all K-finite functions on K.

Step 2. I now apply:

4.11. Lemma. *If $f \neq 0$ is in $C(K)$, then $\sigma(f) \neq 0$ for some irreducible σ).*

Proof. Let $h = f^* \cdot f$. It suffices to prove this for h instead of f. I recall that L_h is a self-adjoint compact operator, and hence that $\mathrm{L}^2(K)$ may be decomposed into at most a countable number of eigenspaces plus the kernel of h. Suppose $L_h v = \mu v$ with $\mu \neq 0$. Then the space of all eigenvectors for μ is stable under right multiplication by G, so according to Lemma 4.2 v is contained in a subspace of K-finite vectors stable under $K \times K$. One can choose σ to be one of its irreducible constituents. ▯

Step 3. Suppose f to be in $C(K)$. I now claim that $\sigma(f) \neq 0$ if and only if f is orthogonal to the space of matrix coeffiicients of σ. This is a consequence of Corollary 3.17.

The proof of Peter-Weyl is concluded. ▯

4.12. Corollary. *In any continuous representation of K the K-finite vectors are dense.*

Proof. I'll prove this first for the space $C(K)$.

Suppose f, φ to be in $C(K)$. Let $h = f * f^*$. Then L_h is a self-adjoint compact operator, say with unit eigenvectors ψ_i, eigenvalues μ_i. We can write

$$\varphi = \sum c_i \psi_i + \Psi$$

with $\sum_i |c_i|^2 < \infty$, Ψ on the kernel of L_f. Choose n so large that $\sum_{i>n} |c_i|^2 < \varepsilon$.

$$L_f\varphi = L_f\Big(\sum\nolimits_{i\le n} c_i\psi_i\Big) + L_f\Big(\sum\nolimits_{i>n} c_i\psi_i\Big)\,.$$

But L_f is a bounded operators from $\mathrm{L}^2(K)$, say with bound M. Therefore the secnd term above is $\le M\varepsilon$. So

$$\Big| L_f\varphi - \Big(\sum\nolimits_{i>n} c_i\psi_i\Big)\Big| < M\varepsilon\,.$$

Therefore, $L_f\varphi$ may be well approximated by K-finite functions. But then φ may be well approximated in turn by convolutions $L_f\varphi$, if we choose f to vary over a Dirac sequence.

Now to deal with the general case. I begin with the observation that a vector v in a representation is K-finite if and only if $v = \sigma(f)u$ for some K-finite function f in $C(X)$, u in V.

Suppose (σ, V) to be a continuous representation of K. Suppose (f_n) to be a Dirac sequence on K. Given a neighbourhood U of 0 in V and v in V, $\sigma(f_n)(v)$ will lie in $v + U/2$ for n large. Fix n for the moment, and choose a sequence (h_i) of K-finite functions in $C(K)$ converging uniformly to f_n. But then $\sigma(h_i)(v)$ converges to $\sigma(f_n)(v)$, and will lie in $v + U$ for large i. ▯

There is a slightly paradoxical flavour to the approximation of a function cin $C(X)$ by K-finite functions, since there are well known examples of functions that are not well approximated by the sequence of its Fourier expansions. When $K = \mathbb{S}$, the classical theorem of Fejér explains some paradox.

4.13. Corollary. *Every irreducible continuous representation of K is finite-dimensional.*

Remarks. Two representations are **equivalent** if they are isomorphic. The exposition of the Peter-Weyl theorem in [Stein:1970] makes a sharp distinction between a representation and an equivalence class of representations. I do not make a sharp distinction, and I'd like to say something about that.

If f is in $C(K)$, I have defined its Fourier transform to be the collection of operators $\widehat{\pi}(f)$ as π varies over the set of irreducible complex representations of K. This does not make much sense, on the face of it. For one thing, the realization of an irreducible representation is by no means canonically defined. The only truly intrinsic object associated to a representation is its trace function on K. The only representation that is canonical is that of $K \times K$ on $V \otimes \widehat{V}$, which is embedded canonically in $C(K)$ in terms of matrix coefficients. Still, I don't see that my sloppy treatment has bad effects.

5. Lie groups

So far, everything I have said is valid for all compact groups. Compact Lie groups are special.

5.1. Corollary. *Any compact Lie group K has a faithful, continuous representation of finite dimension.*

This is certainly not true of all compact groups—for example, $\mathrm{GL}_n(\mathfrak{k})$ with $\mathfrak{k}$ an archimedean local field.

Proof. Let H be the intersection of all the kernels of finite-dimensional representations. It is closed in K. Let $\mathcal{P}_H$ be the projection

$$[\mathcal{P}_H f](k) = \int_H f(hk)\, dh$$

onto the subspace of H-fixed functions in $\mathrm{L}^2(K)$. The group H is normal and closed in K, and the subspace of its invariants closed in $\mathrm{L}^2(K)$.

Suppose $H \neq \{1\}$. By Urysohn's Lemma, there exists on K a non-negative continuous function f such that $f(1) = 0$, $f(h) = 1$ for some h in H. Consequently the open subset of f in $\mathrm{L}^2(K)$ such that $\mathcal{P}_H f - f \neq 0$ is not empty. By the Peter-Weyl theorem, the space $\mathrm{L}^2(K)$ is a direct Hilbert sum of K-stable finite-dimensional subspaces, and these are hence dense in $\mathrm{L}^2(K)$. Therefore there exist K-finite functions that are not constant on H, contradicting the definition of H.

We now know that $H = \{1\}$. Hence for every $k \neq 1$ in K there exists a representation π with $\pi(k) \neq I$. Since K is a Lie group, there exists an open neighbourhood X of 1 which contains no subgroup other than $\{1\}$. Let Y be the complement of X in K. Since Y is compact, there exists a finite set Σ of representations π such that for every k in Y there exists some π in Σ with $k \notin \mathrm{Ker}(\pi)$. The kernel of $\Pi = \bigoplus_\Sigma \pi$ then contains no element of Y, so it must be contained in U. It must therefore be $\{1\}$, so Π is a faithful representation of K. ∎

If K is a compact Lie group, then the Fourier decomposition of a smooth function f looks much like that of a smooth function on the unit circle, but I'll not prove that here. There are two approaches, one concerned with the asymptotics of the eigenvalues of the Laplacian on an arbitrary Riemannian manifold (see Chapter 12 of [Taylor:1981]), the other relying on an explicit classification of the irreducible representations for a connected compact group in terms of characters of a maximal torus, as in [Borel:1972], pages 24–35.

Another thing I'll not prove here is that every compact subgroup of some GL_n is a Lie group. There are several ways to see this, but one very satisfactory approach is to see that it is in fact an algebraic group, defined as the zero set of polynomials (a classic result due, although in a rather different formulation, to Tannaka, and reformulated in [Chevalley:1946]).

6. Smooth representations

In this section, let G be a Lie group.

If (π, V) is a continuous representation of G, a vector in V is called **smooth** if the function $\pi(g)v$ is a smooth function on G with values in V. In this case

$$\lim_{t\to 0} \frac{\pi(\exp(tX))v - v}{t} = \pi(X)v$$

exists for all X in $\mathfrak{g}$.

The representation is said to be smooth if all vectors in V are smooth. If (π, V) is any continuous representation let

$$V^{(m)} = \{v \in V \mid g \mapsto \pi(g)v \ \text{ is in } \ C^m(G, V)\}$$

This space, assigned the norms $\big\|\pi(X)v\big\|_\rho$ for X of order at most m in $U(\mathfrak{g})$, is a quasi-complete Hausdorff TVS, and complete if V is.

6.1. Proposition. *If (π, V) is a continuous representation of G then*

(a) the subspace V^∞ is a continuous representation of G;
(b) for each X in $U(\mathfrak{g})$ the operator $\pi(X)$ takes V^∞ to itself, and is continuous;
(c) for every f in $C_c^m(G)$ and v in V the vector $\pi(f)v$ is differentiable of order m, and

$$\pi(X)\pi(f) = \pi(L_X f)\,.$$

for X in $U^m(\mathfrak{g})$.

Proof. Only the last requires confirmation. We have

$$\begin{aligned} \pi(f) &= \int_G f(x)\pi(x)\,dx \\ \pi\big(\exp(tX)\big)\pi(f) &= \int_G f(x)\pi\big(\exp(tX)x\big)\,dx \\ &= \int_G f\big(\exp(-tX)x\big)\pi(x)\,dx\,. \end{aligned}$$

The map from V to $C(G, V)$ taking v to the function $\pi(g)v$ is injective. The image of V^∞ is a closed subspace of $C^\infty(G, V)$, and its topology is the inherited one.

In Proposition 6.1, m is allowed to be ∞. Choosing a Dirac sequence of smooth functions, we see:

6.2. Corollary. *In any continuous representation of G the smooth vectors are dense.*

6.3. Proposition. *Any continuous finite-dimensional representation of a Lie group is smooth.*

Proof. If (π, V) is continuous, the image of $C_c^\infty(G)$ in $\mathrm{End}(V)$ is closed and by Proposition 2.7 therefore contains I. If $\pi(f) = I$, then $v = \pi(f)v$ for every v, which is therefore smooth. ∎

Remark. If (π, V) is a continuous representation, we have seen that the space spanned by vectors $\pi(f)v$ with f in $C_c^\infty(G)$ is contained in the space of smooth vectors. It is a remarkable result found in [Dixmier-Malliavin:1978] that if V is Fréchet these two spaces are the same.

7. Representations of G and of $(\mathfrak{g}, K)$

In this section, let G be the group of $\mathbb{R}$-rational points on a Zariski-connected reductive group defined over $\mathbb{R}$, K a maximal compact subgroup of G. According to Cartan's fixed point theorem, all choices of K are conjugate to each other by an element of the connected component of G. Suppose A to be the group of $\mathbb{R}$-rational points on a maximal split torus of G. According to §14 of [Borel-Tits:1972], A meets all connected components of G. This implies the group K meets all of the connected components of G. (This is a special case of an older result of [Mostow:1955] about arbitrary Lie groups with a finite number of connected components.) The group G acts trivially by conjugation on $Z(\mathfrak{g})$, the center of the enveloping algebra of G.

A continuous representation of G is **admissible** if the dimension of each K-isotypic component has finite dimension. Admissible representations are ubiquitous as well as important. All irreducible unitary representations of G are known to be admissible, and admissible representations are those occurring in various decompositions of natural representations of G, such as that of $G \times G$ on $C_c^\infty(G)$ or $\mathrm{L}^2(G)$. In short, they are basic objects in the theory.

If (π, V) is any continuous representation of G, let $V_{(K)}$ be the subspace of K-finite vectors—that is to say, contained in a K-stable finite-dimensional subspace of V.

7.1. Proposition. *If $(\pi, V$ is an admissible continuous representation of G, then for every v in $V_{(K)}$ there exists a function f in $C_c^\infty(G)$ such that $\pi(f)v = v$.*

Proof. If ξ is an idempotent in $C(K)$ such that $\pi(\xi)$ is the identity on the finite dimensional K-isotypic component U, then the image of $\pi(\xi)\pi(C_c^\infty(G))\pi(\xi)$ is a closed subalgebra of the finite-dimensional algebra $\mathrm{End}_{\mathbb{C}}(U)$. But if we choose a Dirac sequence φ_i in $C_c^\infty(G)$, then $\pi(\xi)\varphi_i\pi(\xi)$ will converge to the identity of U. So the identity operator is in the image, too. ∎

7.2. Corollary. *If V is an admissible representation of G, then the subspace $V_{(K)}$ is contained in V^∞ and is stable with respect to $\mathfrak{g}$.*

Proof. The first claim follows immediately from the Proposition and Proposition 6.1. If U is a finite-dimensional subspace, the map $X \otimes u \mapsto \pi(X)u$ is a surjection of K-spaces $\mathfrak{g} \otimes U \to \pi(\mathfrak{g})U$. ∎

7.3. Corollary. *Every K-finite vector in an admissible representation (π, V) of G is also $Z(\mathfrak{g})$-finite. If π is irreducible, there exists a homomorphism from $Z(\mathfrak{g})$ to $\mathbb{C}$ through which elements of $Z(\mathfrak{g})$ act.*

This homomorphism is traditionally called (for reasons that escape me) the **infinitesimal character** of π.

Many technical problems are to be found in the theory of infinite-dimensional representations of G that don't exist for finite-dimensional ones. The most serious arise because many different continuous representations of G are all in some sense equivalent, even though analytically of a very different nature. Some are quite complicated, in fact unnecessarily complicated. The sort of thing I have in mind is illustrated by representations of $\mathrm{SL}_2(\mathbb{R})$ associated to its action on $\mathbb{P} = \mathbb{P}^1(\mathbb{R})$. The spaces of analytic functions, smooth functions, and locally L^2 functions on $\mathbb{P}$ are for most purposes best considered as different incarnations of the same beast. What they all have in common is the same underlying space of K-finite functions. This, happily, is a representation of the Lie algebra $\mathfrak{sl}_2$, if not of the group itself.

Suppose given a representation of $\mathfrak{g}$ and a continuous representation of K simultaneously on a space V. I'll denote noth representations by π. It is called an admissible representation of the pair $(\mathfrak{g}, K)$ if

(a) as a representation of K it is an algebraic direct sum of irreducible finite-dimensional representations of K, each with finite multiplicity;
(b) the two representations of $\mathfrak{k}$, as Lie algebra of K and as subalgebra of $\mathfrak{g}$, are the same.

These are often called **Harish-Chandra modules**. If V is an admissible representation of G then $V_{(K)}$ is an admissible representation of $(\mathfrak{g}, K)$.

In practice, in the theory of representations of the reductive group G, one works with admissible representations of $(\mathfrak{g}, K)$ rather than continuous representations of G. We shall see in a moment how this can be justified. The fundamental question is this: *to what extent does a representation of the Lie algebra determine that of the group?* One problem that the representation of the Lie algebra alone can't handle at all is the behaviour of the representation on connected components of G other than that of the identity, but that problem is addressed by including the action of K, which meets all connected components. (This is true for arbitrary Lie groups with a finite number of connected components—see [Mostow:1955].)

Before taking up the justification of the definition of Harish-Chandra modules, let me point out why replacing the group G by the pair $(\mathfrak{g}, K)$ is a good idea. Suppose for the moment that $G = \mathrm{SL}_2(\mathbb{R})$, and let $\mathbb{P} = \mathbb{P}^1(\mathbb{R})$. The group G acts on the right this space through the action

$$v = \begin{bmatrix} x \\ y \end{bmatrix} \longmapsto g^{-1} v \,.$$

It therefore acts on the left on various spaces of functions on $\mathbb{P}$. Among these are the space $C(\mathbb{P})$ of continuous functions, the space $C^\infty(\mathbb{P})$ of smooth functions, the space $C^\omega(\mathbb{P})$ of real analytic functions, and $\mathrm{L}^{2,\mathrm{loc}}(\mathbb{P})$ of locally square-integrable functions. All of them contain exactly the same space of K-finite functions, which as a representation of K is the direct sum of characters ε^{2n} of K where

$$\varepsilon\colon \begin{bmatrix} c & -s \\ s & c \end{bmatrix} \longmapsto c + is \,.$$

As you can imagine, a great deal of simplification takes place if one looks just at this last space—analysis will be replaced by algebra.

There are a number of facts that justify the switch from representations of G to those of $(\mathfrak{g}, K)$. I state them here briefly:

(a) every Harish-Chandra module is the restriction to K-finite vectors of some continuous representation of G;
(b) if (π, V) is an admissible representation of G, the map taking $U \subseteq V_{(K)}$ to its closure in V is a bijection between $(\mathfrak{g}, K)$-stable subspaces of $V_{(K)}$ and closed G-stable subspaces of V;
(c) the map taking (π, V) to $V_{(K)}$ is a bijection between irreducible unitary representations of G and irreducible Harish-Chandra modules with a positive definite metric invariant with respect to $(\mathfrak{g}, K)$;

(d) there exist exact functorial assignments of smooth representations of G to Harish-Chandra modules, inverse to the restriction map.

I'll look at (a) in the next section. As for the rest, I'll just look at them briefly without going into detail.

I'll sketch the proof of (b) in a moment.

[Borel:1972] proves (c).

There are several ways to assign continuous representations of G to Harish-Chandra modules in a functorial manner. One is in a sense minimal, and is discussed in [Kashiwara-Schmid:1994]. The other perhaps the most natural, in that it incorporates features most useful in applications. The original ideas were in joint work by Nolan Wallach and myself. This is explained in [Casselman:1990]. A different approach to the same construction is to be found in [Bernstein-Krötz:2010].

I now sketch the proof of (b).

7.4. Proposition. *If Φ is a distribution on G which is left- or right-K-finite as well as $Z(\mathfrak{g})$-finite, then it is a real-analytic function on G.*

Proof. Let π be the left or right regular representation of G, depending on the assumption regarding K.

Since Φ is the sum of its K-isotypic components, one may as well assume that Φ is itself an isotypic component, hence that $\pi(C_K)\Phi = \lambda\Phi$ for some λ, where C_K is the Casimir element of $U(\mathfrak{k})$. Since C_K commutes with all of $Z(\mathfrak{g})$, $\pi(XC_K)\Phi = \pi(C_K)\pi(X)\Phi = \lambda\pi(X)\Phi$ for all X in $Z(\mathfrak{g})$. Since Φ is $Z(\mathfrak{g})$-finite

$$\prod_1^n \big(\pi(C) - \mu_i\big)\Phi = 0$$

for some set of scalars μ_i, where now C is the Casimir element of $U(\mathfrak{g})$. Let $\omega = C + 2C_K$, and let

$$\Phi_k = \prod_1^k \big(\pi(C) - \mu_i\big)\Phi \quad (0 \le k \le n)$$

and in particular $\Phi_0 = \Phi$. Thus

$$\begin{aligned}\big(\pi(C) - \mu_k\big)\Phi_{k-1} &= \Phi_k \\ \big(\pi(C) - \mu_n\big)\Phi_{n-1} &= 0\end{aligned}$$

leading to

$$\begin{aligned}\big(\pi(\omega) - (2\lambda + \mu_k)\big)\Phi_{k-1} &= \Phi_k \\ \big(\pi(\omega) - (2\lambda + \mu_n)\big)\Phi_{n-1} &= 0\,.\end{aligned}$$

Since ω is an elliptic operator on G with analytic coefficients, each Φ_k is analytic. ▯

7.5. Corollary. *If (π, V) is an admissible representation of G, v a K-finite vector and $\widehat{v}$ in $\widehat{V}$ any continuous linear functional on V, then the matrix coefficient $\langle \pi(g)v, \widehat{v}\rangle$ is analytic.*

7.6. Theorem. *If $(\pi, \overline{V})$ is admissible, the map taking W to $\overline{W}$ is a bijection of $(\mathfrak{g}, K)$-stable subspaces of V and closed G stable subspaces of $\overline{V}$.*

Proof. The difficult point, and the one that convinces most strongly that representations of $(\mathfrak{g}, K)$ are a reasonable thing to serve as formal substitutes for representations of G, is the claim that if W is a $(\mathfrak{g}, K)$-stable subspace then its closure $\overline{W}$ is G-stable. By the Hahn-Banach theorem, in order to show that $\overline{W}$ is G-stable, it suffices to show that if F is a linear function on V that vanishes on W, then $F\big(\pi(g)w\big) = 0$ for all w in W. Because K meets all components of G and W is K-stable, it suffices to show this for g in the connected component of G. But according to Corollary 7.5 this is an analytic function of g, and therefore it suffices to show that the coefficients of its Taylor series at 1 vanish. However, these coefficients are determined by the constants $F(\pi(X)w)$, which vanish by assumption. ∎

8. Realization

I call a Harish-Chandra module (π, v) **realizable** if it is the representation of $(\mathfrak{g}, K)$ on the K-finite vectors in admissible continuous representation of G. As I have said in the previous section, all Harish-Chandra modules are realizable. I also mentioned there that there is in fact a canonical way to realize every Harish-Chandra module, but that seems not to be so important as the simple existence of *some* realization.

The original result about realizability was one by Harish-Chandra, which asserted that ever irreducible admissible representation of$(\mathfrak{g}, K)$ could be realized as subquotient of some principal series. A later and more interesting proof of this was found by [Beilinson-Bernstein:1982]. I myself might have been the first the prove that every finitely generated $(\mathfrak{g}, K)$-module was realizable. This proof showed that formal solutions of the differential equations satisfied by matrix coefficients were in fact converging solutions, but this proof was never published. Other proofs have been found since. The existence of matrix coefficients is presumably the most important fact about realizable representations,a because the asymptotic behaviour of matrix coefficients is an extremely important part of the theory. Jacquet and Langlands essentially postulated this in the their book on GL_2, where it is disguised in terms of an action of the Hecke algebra.

9. *Appendix.* Tensors and homomorphisms

My goal in this appendix is to recall, and comment on, a number of conventions in linear algebra that are constantly called on in representation theory, but do not depend on special properties of G or V. The novel point is to explain what might seem at first somewhat arbitrary choices, and how they relate to each other.

Throughout, G will be a group and $\mathfrak{k}$ a field of characteristic 0. Representations of G will be on vector spaces over $\mathfrak{k}$. If V is a vector space over $\mathfrak{k}$, $\widehat{V}$ will be a subspace of its $\mathfrak{k}$-linear dual. The space $C(G) = C(G, \mathfrak{k})$ will be that of functions on G with values in $\mathfrak{k}$. If G is locally compact, representations will be assumed to be continuous, functions in $C(G)$ will be continuous, and $C_c(G)$ will be the subspace of functions with compact support.

DUALITY. Suppose (π_1, V_1) and (π_2, V_2) to be two representations of G. There are two natural ways to define a representation of $G \times G$ on the vector space $\mathrm{Hom}_{\mathbb{C}}(V_1, V_2)$, the difference being a matter of order:

$$[\mathrm{Hom}(\pi_1, \pi_2)](g_1, g_2)\colon \quad F \longmapsto \pi_2(g_1)\, F\, \pi_1(g_2^{-1})$$
$$[\mathfrak{Hom}(\pi_1, \pi_2)](g_1, g_2)\colon \quad F \longmapsto \pi_2(g_2)\, F\, \pi_1(g_1^{-1})$$

My default choice will be the first, even though the mixing of indices is awkward. This choice is easily motivated. The group $G \times G$ acts on G, hence also $C(G)$, from right and left, and given a representation $\pi, V)$ of G, it is reasonable to require compatibility with the action on $\mathrm{End}(V)$. With my choice, the map $f \mapsto \pi(f)$ is $G \times G$-equivariant:

$$\begin{aligned}\pi(L_{g_1} R_{g_2} f) &= \int_G f(g_1^{-1} x g_2) \pi(x)\, dx \\ &= \int_G f(x) \pi(g_1 x g_2^{-1})\, dx \\ &= \pi(g_1)\pi(f)\pi(g_2^{-1})\,.\end{aligned}$$

The two representations agree on the diagonal copy of G in $G \times G$. In particular, for both the space $\mathrm{Hom}_G(V_1, V_2)$ of G-equivariant linear maps from V_1 to V_2 is the subspace of invariants with respect to this diagonal copy.

A special case of this is the dual representation of G on $\widehat{V}$, for which

$$\langle v, \widehat{\pi}(g)\widehat{v}\rangle = \langle \pi(g^{-1})v, \widehat{v}\rangle\,.$$

This raises a question. Suppose (π, V) to be a representation of G, $\widetilde{V}$ its dual. The problem that arises here is how to write the pairing of V and $\widetilde{V}$: $\langle v, \widetilde{v}\rangle$ or $\langle \widetilde{v}, v\rangle$? My choice here is the first, because it is consistent with my choice regarding homomorphisms.

The problem with my choice is that it is not compatible with the standard matrix convention. According to this, vectors are columns and linear functions are rows. If $\widehat{v}$ is a row and v is a column, the pairing is the matrix product $\widehat{v}{\cdot}v$, whereas I write it $\langle v, \widehat{v}\rangle$, alas. The ultimate source of troubles like this seems to be how we apply maps: the value of the function F applied to x is $F(x)$. But the value at x of the composite

$$U \xrightarrow{F} V \xrightarrow{G} W$$

is then $G(F(x))$. There have been attempts to change these conventions (I seem to recall one in the book from which I learned 'modern' algebra), but have not succeeded.

I have made a default choice between the two, but even so the other is inevitable. The map taking f to its contragredient $\widehat{f}$ is from $\mathrm{Hom}(V_1, V_2)$ to $\mathrm{Hom}(\widehat{V}_2, \widehat{V}_1)$. But it is not equivariant with respect to my choice. Instead, it is equivariant if we choose the representation Hom on the first and $\mathfrak{Hom}$ on the second. This is because $f \mapsto \widehat{f}$ is contravariant.

Let $\mathfrak{End}(V)$ be the ring of all operators from V to itself of finite rank. If $\widehat{v}$ is in $\widehat{V}$ and v in V, the operator

$$\varepsilon_{v\otimes\hat{v}}\colon\ u \longmapsto \langle u, \widehat{v}\rangle\, v$$

is in $\mathfrak{End}(V)$, with trace $\langle v, \widehat{v}\rangle$. Its image consist of the line through v.

9.1. Lemma. *The map taking $v \otimes \widehat{v}$ to $\varepsilon_{v\otimes\widehat{v}}$ is a $G \times G$-equivariant bijection of $V \otimes \widehat{V}$ with $\mathfrak{End}(V)$.*

Proof. Let f be in $\mathfrak{End}(V)$, with image contained in the finite-dimensional space U. Let (u_i) be a basis of U. Then $f(v) - \sum f_i(v) u_i$ for some uniquely defined linear maps f_i in $\widehat{V}$.

There is another place in which $\mathfrak{Hom}$ appears naturally. If Φ is in $\mathrm{End}(V)$ and F in $\mathfrak{End}(V)$, the product $\Phi\, F$ is of finite rank, and the pairing

$$\mathrm{trace}\,(\Phi\, F)$$

identifies $\mathrm{End}(V)$ with a subspace of the linear dual of $\mathfrak{End}(V)$. The pairing is not quite equivariant, however, with respect to my 'default' G-actions. It is, however, if $G \times G$ acts on $\mathrm{End}(\pi)$ by my choice, and on $\mathfrak{End}(\pi)$ by the other.

FUNCTIONS. Given v in V, $\widehat{v}$ in $\widehat{V}$, the associated **matrix coefficient** is the function

$$\mu_{v \otimes \widehat{v}} \colon g \longmapsto \langle \pi(g^{-1})v, \widehat{v} \rangle$$

in $C(G)$.

9.2. Proposition. *The map taking $v \otimes \widehat{v}$ to $\mu_{v \otimes \widehat{v}}$ is a $G \times G$-equivariant map from $V \otimes \widehat{V}$ to $C(G)$.*

If π is irreducible, then according to Proposition 3.7 this is an embedding.

$$\mu_{v \otimes \hat{v}}(g) = \langle v, \widehat{\pi}(g)\widehat{v} \rangle = \langle \pi(g^{-1})v, \widehat{v} \rangle .$$

There is another way to define certain matrix coefficients. Define a map τ from $\mathfrak{End}_{\mathfrak{k}}(V)$ to $C(G)$:

$$\tau \colon f \longmapsto \left[g \longmapsto \mathrm{trace}\left(\pi(g^{-1}) f \right) \right] .$$

9.3. Proposition. *The diagram*

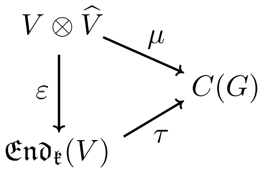

is a commutative diagram of $G \times G$-representations.

One has to be careful about the relationship between tensors and linear maps when V has infinite dimension. One has to be especially careful for topological vector spaces. Only for nuclear vector spaces is there a reasonable identification of a (topological) tensor product with a space of linear maps.

10. References

1. Armand Borel, **Représentations de groupes localement compacts**, *Lecture Notes in Mathematics* **276**, Springer, 1972.

2. Nicholas Bourbaki, **Integration I**, Springer. Translated from the original French edition, which is dated 1959–1965.

3. ——, **Topological vector spaces**, Springer, translated from the French original dated 1981.

4. Jacques Dixmier and Paul Malliavin, 'Factorisations de fonctions et de vecteurs indéfiniment différentiables', *Bulletin des Sciences Mathématiques* **102** (1978), 305–330. 1978

5. Claude Chevalley, **Theory of Lie Groups**, Princeton University Press, 1946.

6. Edwin Hewitt and Kenneth Ross, **Abstract harmonic analysis**, volume II, Springer, 1969.

7. Daniel Mostow, 'Self-adjoint groups', *Annals of Mathematics* **62** (1955), 44–55.

8. Eli Stein, **Topics in harmonic analysis**, Princeton University Press, 1970.

9. André Weil, **L'intégration dans les groupes topologiques**, Hermann, 1965.

10. Wikipedia on Fejér's Theorem:

https://en.wikipedia.org/wiki/Fejers_theorem