Harmonic Analysis and Representation Theory for Semisimple Lie Groups

Matthew Dawson

February 24, 2015

Contents

Chapter 1

Motivation and Basic Tools

1.1 What is a group representation?

- A representation of a group is nothing more or less than an action of a group on a vector space.
- In other words, it is a way to realize a particular group as a collection of linear operators on a vector space.
- They appear in physics, because quantum mechanics sayst that the state of a physical system should be identified with a vector in a Hilbert space. Thus, any group of transformations to the system should correspond to a group of linear operators.
- They appear in the theory of linear PDEs, because the study of linear PDEs is basically equivalent to studying the spectral decomopsition of the linear operator which gives the equation. Sometimes, the equation possesses certain symmetries (i.e., the linear operator is invariant under a certain group of transformations). In that case, the solution spaces to the equation correspond to group representations.
- They appear in probability theory, through the Langlands program and automorphic forms, which pertain to harmonic analysis over the quotients of semisimple Lie groups by certain discrete subgroups.

1.2 The Definition of Group Representations

We begin with the definition of a representation of a group on a vector space:

Definition 1.2.1. *A* ***representation*** *of a group G is a pair (π, V), where V is a vector space over a field $\mathbb{F}$ and π is a homomorphism*

$$\pi : G \to \mathrm{GL}(V),$$

where as usual $GL(V)$ denotes the group of invertible linear operators from V to V.

In other words, a representation is just an action of a group G in the category of vector spaces over some field $\mathbb{F}$.

For finite groups, this definition is quite sufficient. However, when we consider a continuous group like $\mathbb{R}$, talking about *all* representations is basically as difficult as talking about all *functions* on $\mathbb{R}$. What we really care about are *continuous* functions on $\mathbb{R}$ and *continuous* representations of $\mathbb{R}$. Of course, this requires a notion of *topological groups* and *topological vector spaces*:

Definition 1.2.2. *A topological group is a group G with a topology such that the maps*

$$\begin{array}{ccc} G \times G & \to & G \\ (g,h) & \mapsto & gh \end{array} \quad \textit{and} \quad \begin{array}{ccc} G & \to & G \\ g & \mapsto & g^{-1} \end{array}$$

are continuous.

A vector space V over $\mathbb{F}$ (where $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$) is a topological vector space if the maps

$$\begin{array}{ccc} V \times V & \to & V \\ (v,w) & \mapsto & v+w \end{array} \quad \textit{and} \quad \begin{array}{ccc} \mathbb{F} \times V & \to & G \\ (\lambda,v) & \mapsto & \lambda v \end{array}$$

are continuous.

We will always assume that our topological groups are Hausdorff, locally compact, and second-countable (and hence separable). On the other hand, our topological vector spaces will always be Hausdorff, complete, and locally convex (that is, every neighborhood of a point contains a smaller neighborhood of that point which is convex and balanced). However, we will have occasion to discus group representations on a locally convex topological vector space which is not even first-countable.

So a **representation** of a topological group G on a topological vector space V (over $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$) *should* be a continuous homormoprhism

$$\pi : G \to \mathrm{GL}(V),$$

where $\mathrm{GL}(V)$ now denotes the group of continuous linear maps from V to V with continuous inverses.

For a finite-dimensional vector space V, there is only one useful topology we can put on $\mathrm{GL}(V)$. But there are many choices of topologies which we may put on $\mathrm{GL}(V)$ if V is infinite-dimensional. Which one should we use here? If V is a Banach space, then $\mathrm{GL}(V)$ may be turned into a Banach space under the operator norm. Or we could give $\mathrm{GL}(V)$ the strong operator topology or even the weak operator topology.

The answer to which topology is "correct" for defining continuous representations comes down to practical experience: the representations which are most important to us tend to be continuous in certain topologies and discontinuous in others. We will see this in a moment when we begin to talk about regular representations, but for now, it suffices to say that we will always use the strong operator topology on $\mathrm{GL}(V)$. (That is the topology under which a sequence $\{T_n\}_{n\in\mathbb{N}} \subseteq \mathrm{GL}(V)$ converges to an operator $T \in \mathrm{GL}(V)$ if and only if

$$\lim_{n\to\infty} T_n v = Tv$$

for all $v \in V$.) Thus, the "correct" definition of a representation of a topological group is:

Definition 1.2.3. *A **strongly continuous representation** of a topological group G on a topological vector space V over $\mathbb{R}$ or $\mathbb{C}$ is a homomorphism $\pi : G \to \mathrm{GL}(V)$ such that the maps*

$$\begin{array}{ccc} G & \to & V \\ g & \mapsto & \pi(g)v \end{array}$$

are continuous for each $v \in V$.

We will always assume that our group representations are strongly continuous.

Note that a strongly continuous representation $\pi : G \to \mathrm{GL}(V)$ has the property that

$$\begin{array}{ccc} G \times V & \to & V \\ (g, v) & \mapsto & \pi(g)v \end{array}$$

is separately continuous—that is, continuous in each parameter when the other is fixed (continuity of the maps $v \mapsto \pi(g)v$ for each $g \in G$ follows from the fact that $\pi(g) : V \to V$ is continuous). However, strongly continuous representations usually satisfy a stronger continuity condition:

Theorem 1.2.4. *A representation π of a topological group G on a Frechét space V is strongly continuous if and only if the map*

$$\begin{array}{ccc} G \times V & \to & V \\ (g, v) & \mapsto & \pi(g)v \end{array}$$

is jointly continuous (i.e., continuous as a map from $G \times V$ to V).

Exercise 1. Prove Theorem 1.2.4. *Hint:* Use the Banach-Steinhaus Theorem (also known as the Uniform Boundedness Principle), which holds true for all Frechét spaces.

Suppose that (π, V) is a strongly continuous representation of a group G on a complete topological vector space V over $\mathbb{R}$ or $\mathbb{C}$. Now suppose that there is a closed subspace $W \subseteq V$ such that $\pi(g)W \subseteq W$ for all $g \in G$. Then we say that W is a **closed invariant subspace** of V and notice that we can define a new continuous representation π^W of G on W by setting $\pi^W(g)w = \pi(g)w$ for each $g \in G$ and $w \in W$. The fact that π^W is a strongly continuous representation of G then follows from the fact that W is a *closed* subspace of V. We say that π^W is a **subrepresentation** of π.

As in any good mathematical category, we can use this notion of *subrepresentations* to define *quotient representations*. Suppose that (π, V) is a strongly continuous representation of G with a subrepresentation (π^W, W) is a subrepresentation. The fact that W is a closed subspace of V means that V/W is again a topological vector space. We can define the **quotient representation** π/π^W on V/W by setting:

$$(\pi/\pi^W)(g)(v+W) = \pi(g)v + W.$$

We leave it to the reader to show that this representation is well-defined and continuous.

A representation (π, V) of G is **irreducible** if V possesses no closed, invariant subspaces besides V and $\{0\}$. Irreducible representations should be thought of as the "basic building blocks" of representations.

To see why irreducible representations are important, we assume for now that V is finite-dimensional. Then V possesses a maximal invariant subspace $V_1 \subsetneq V$. It follows that the quotient representation $(\pi/\pi^{V_1}, V/V_1)$ is irreducible [prove this! Hint: any invariant subspace of V/V_1 corresponds to an invariant subspace of V which contains V_1]. if $V_1 \neq \{0\}$, then we choose a maximal invariant subspace $V_2 \subsetneq V_1$ and note that $(\pi^{V_2}/\pi^{V_2}, V_1/V_2)$ is irreducible. Because V is finite-dimensional, we can continue this process until $V_i = \{0\}$ for some i. We thus arrive at the so-called **composition series** for (π, V):

$$\{0\} = V_i \subsetneq V_{i-1} \subsetneq V_2 \subsetneq V_1$$

where each V_k is an invariant subspace of V such that $(\pi^{V_k}/\pi^{V_{k+1}}, V_k/V_{k+1})$ is an irreducible representation of G. We made several choices of V_k's throughout the process, but one can show that the length of the composition series will always be the same, and that we will always get the same quotient representations of G, up to a permutation on the order.

Speaking of representations being "the same," what do we mean be that? We need a notion of morphisms and isomomorphisms for representations, to borrow category theory language. Those will be called intertwining operators and equivalences of representations, respectively:

Definition 1.2.5. *Suppose that (π, V) and (σ, W) are both continuous representations of a group G. Then we say that a continuous linear map $L : V \to W$ is an **intertwining operator** if*

$$L\pi(g) = \sigma(g)L$$

for all $g \in G$. The vector space of all intertwining operators from V to W is denoted by $\mathrm{Hom}_G(V, W)$ or $\mathrm{Hom}_G(\pi, \sigma)$, depending on whether we want to emphasize the vector space or the representation.

*We say that (π, V) and (σ, W) are **equivalent** if there is an intertwining operator $L : V \to W$ which is continuously invertible, so that $L^{-1} : W \to V$ is an intertwining operator. We write $\pi \cong \sigma$ or $V \cong_G W$ if (π, V) and (σ, W) are equivalent.*

Exercise 2. Suppose that (π, V) and (σ, W) are both continuous representations of a group G. Show that $\mathrm{Hom}_G(\pi, \sigma)$ is a closed subspace of $\mathcal{B}(V, W)$ under the strong operator topology. (Here $\mathcal{B}(V, W)$ is the space of all continuous operators from V to W.)

If V and W are Banach spaces, then this immediately implies closedness of $\mathrm{Hom}_G(\pi, \sigma)$ in the operator norm topology on $\mathcal{B}(V, W)$. As another consequence, this implies that the space $\mathrm{Hom}_G(\pi, \pi)$ is a closed subalgebra of the algebra $\mathcal{B}(V)$ of lienar maps from V to V (under the strong operator topology, of course). These two facts are very important for the connection between unitary representations of groups and representations of C^*-algebras.

1.3 Unitary Representations and Schur's Lemma

To motivate the definition of unitary representations, we begin with a simple method for constructing new representations from old ones. Suppose that (π, V) and (σ, W) are continuous representations of G. We form the topological vector space $V \oplus W$ and now define a new representation, called $\pi \oplus \sigma$, on $V \oplus W$ as follows:

$$\pi \oplus \sigma(g)(v, w) = (\pi(v), \sigma(w)).$$

We leave it to the reader to show that π is continuous.

Now suppose that (π, V) is a continuous representation of G and that (σ, W) is a subrepresentation. We have seen how σ and π/σ can provide information about π, but it would be very nice if we could write $V = W \oplus U$, where U is another closed invariant subspace of V. It would then follow that

$$\pi \cong \sigma \oplus \rho,$$

where (ρ, U) is a subrepresentation of (π, V). Unfortunately, it is not always possible to find such a *complement* U for W.

However, if we assume that V is a *Hilbert space* and that π acts by unitary operators on V, then such a complement can always be found.

Definition 1.3.1. *A* ***unitary representation*** $(\pi, \mathcal{H})$ *of a group G on a Hilbert space $\mathcal{H}$ is a strongly continuous homomorphism*

$$\pi : G \to \mathrm{U}(\mathcal{H}),$$

where $\mathrm{U}(\mathcal{H})$ *is the group of unitary operators on the Hilbert space $\mathcal{H}$.*

Lemma 1.3.2. *If $(\pi, \mathcal{H})$ is a unitary representation of a group G and W is a closed invariant subspace of $\mathcal{H}$, then $W^\perp$ is also a closed invariant subspace and $\mathcal{H} \cong_G W \oplus W^\perp$.*

Proof. Recall that

$$W^\perp = \{v \in \mathcal{H} | \langle v, w \rangle = 0 \text{ for all } w \in W\}.$$

Now suppose that $v \in W^\perp$. We want to show that $W^\perp$ is an invariant subspace of $\mathcal{H}$, and so we need to show that $\pi(g)v \in W^\perp$ for all $g \in G$. Now, for all $w \in W$, we see that $\langle \pi(g)v, w\rangle = \langle v, \pi(g)^{-1}w\rangle$ because $\pi(g)$ is a unitary operator. It then follows that $\langle v, \pi(g)^{-1}w\rangle = 0$ because $\pi(g)^{-1}w = \pi(g^{-1})w \in W$ (here we use the fact that W is an invariant subspace of $\mathcal{H}$). Hence $\pi(g)v \in W^\perp$ and so we are done. $\square$

In other words, the advantage of unitary representations is that if they are not irreducible, then they can always be decomposed into direct sums of subrepresentations. In fact, we get the following corollary:

Corollary 1.3.3. *If (π, V) is a unitary representation of G, where V is a finite-dimensional Hilbert space, then π may be decomposed into a direct sum of irreducible subrepresentations.*

From the standpoint of category theory, the category of unitary representations of a group G is distinct from the category of representaions of G over Hilbert spaces. Thus, in order for two unitary representations $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{K})$ to be equivalent, we should require that there be a *unitary* operator $L : \mathcal{H} \to \mathcal{K}$ such that $L\pi(g) = \sigma(g)L$ for all $g \in G$. A priori, we might imagine it possible that two unitary representations could possess in *invertible* intertwining operator $\mathcal{H} \to \mathcal{K}$ but not a *unitary* intertwining operator $\mathcal{H} \to \mathcal{K}$. Then π and σ would be equivalent as *representations* but not as *unitary representations* of G. Fortunately, this scenario can never occur:

Lemma 1.3.4. *Suppose that $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{K})$ are unitary representations and that $L : \mathcal{H} \to \mathcal{K}$ is an invertible intertwining operator. Then there is a unitary intertwining operator $U : \mathcal{H} \to \mathcal{K}$*

Exercise 3. Prove Lemma 1.3.4. *Hint:* Note that $\mathcal{H}$ and $\mathcal{K}$ must be isomorphic as Hilbert spaces. Thus, it suffices to consider the simpler case where $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{H})$ are representations of G on the same space $\mathcal{H}$. We then have an invertible intertwining operator $L \in \mathrm{GL}(\mathcal{H})$. Now use the polar decomposition of L as $L = PU$, where $P = \sqrt{L^*L}$ is positive-definite and $U = P^{-1}L$ is unitary. Show that U is a unitary intertwining operator.

We now give two examples of unitary representations. Thus far, everything we have said has applied equally to representations on real vector spaces and complex vector spaces. However, these two examples will exemplify fundamental difference in behavior between real and complex group representations. This difference in behavior is the reason for which we will discuss only representations on complex vector spaces from here on out.

Example 1.3.5. Consider the group $\mathbb{R}$ of real numbers under addition, as well as the real vector space $\mathbb{R}^2$ under the Euclidean inner product. We define a representation $\pi : \mathbb{R} \to \mathrm{U}(\mathcal{H})$ by setting

$$\pi(x) = \begin{bmatrix} \cos(x) & -\sin(x) \\ \sin(x) & \cos(x) \end{bmatrix}$$

for each $x \in \mathbb{R}$. Then π rotates the Euclidean plane $\mathbb{R}^2$ through an angle of x radians. Since π acts by isometries, it is clear that π acts by orthogonal (i.e., real unitary) operators on $\mathbb{R}^2$. Furthermore, this representation is irreducible, because every nontrivial subspace of $\mathbb{R}^2$ is a line through the origin, and no such line is left invariant under rotations.

Example 1.3.6. As another example, we once again consider the group $\mathbb{R}$ of real numbers, but this time we consider the *complex* vector space $\mathbb{C}$. Then we define a unitary representation σ on $\mathbb{C}$ by defining

$$\sigma(x)z = e^{ix}z$$

for all $x \in \mathbb{R}$ and $z \in \mathbb{C}$.

Of course, $\mathbb{C}$ can also be thought of as a *real* vector space if we forget the complex structure. Then $\mathbb{C}$ can be identified with $\mathbb{R}^2$, and we see that σ and the representation π from the previous exercise are equivalent as representations of G on real vector spaces.

The previous two examples demonstrated that an irreducible representation of $\mathbb{R}$ on a *real* vector space need not be one-dimensional. However, we will show, using *Schur's Lemma*, that every irreducible representation of $\mathbb{R}$ on a complex vector space **must** be one-dimensional.

Schur's Lemma is an incredibly simple theorem with a relatively easy proof which is nevertheless incredibly important to studying unitary representations. We will use it countless times, and so will you if you do work in representation theory.

Theorem 1.3.7 (Schur's Lemma)**.** *Schur's lemma has two parts:*

1. *If (π, V) and (σ, W) are inequivalent irreducible representations of G over real or complex vector spaces V and W, then* $\mathrm{Hom}_G(\pi, \sigma) = \{0\}$.

2. *If π is a representation of G on a Hilbert space $\mathcal{H}$ over $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$ and*

$$\mathrm{Hom}_G(\pi, \pi) = \{\lambda \mathrm{Id} \mid \lambda \in \mathbb{F}\},$$

   then π is irreducible.

3. *If π is an irreducible unitary representation of G over a complex Hilbert space $\mathcal{H}$, then*

$$\mathrm{Hom}_G(\pi, \pi) = \{\lambda \mathrm{Id} \mid \lambda \in \mathbb{C}\}.$$

Proof. Part (1) of the theorem follows quickly from the fact that for any intertwining operator $T : V \to W$, the closure $\overline{T(V)}$ of its image must be a closed invariant subspace of W.

For part (2), we suppose that π is a representation of G on a Hilbert space $\mathcal{H}$ over $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$ and that

$$\mathrm{Hom}_G(\pi, \pi) = \{\lambda \mathrm{Id} \mid \lambda \in \mathbb{F}\}.$$

Now if W is a closed invariant subspace of $\mathcal{H}$, then the orthogonal projection operator $P : \mathcal{H} \to W$ is an intertwining operator in $\mathrm{Hom}_G(\pi, \pi)$. Since P must be a multiple of the identity operator, it follows that W must be equal to $\mathcal{H}$ or $\{0\}$.

For part (3), let π be an irreducible unitary representation of G over a complex Hilbert space $\mathcal{H}$ and suppose that $T \in \mathcal{B}(\mathcal{H})$ is an intertwining operator. Then $\pi(g)^* = \pi(g^{-1})$ for all $g \in G$ and so one quickly shows that T^* is also an intertwining operator. Then $T = \frac{T+T^*}{2} + i\frac{T-T^*}{2i}$, where both $\frac{T+T^*}{2}$ and $\frac{T-T^*}{2i}$ are self-adjoint operators. It thus suffices to consider the case where T is self-adjoint.

Now suppose that T is self-adjoint. By the spectral theorem, there is a projection-valued measure E on $\mathbb{C}$ such that

$$T = \int_{\mathbb{C}} \lambda E_\lambda \, d\lambda.$$

Because $\pi(g)T = T\pi(g)$, it follows that each of the spectral projection operators E_A, where $A \subseteq \mathbb{C}$, also commute with $\pi(g)$.

But the image of any self-adjoint projection operator is a closed invariant subspace of $\mathcal{H}$, so the irreducibility of $\mathbb{C}$ implies that E_A is equal to either 0 or Id for all Borel sets $A \subseteq \mathbb{C}$. It follows that the spectral measure is concentrated on one point λ in $\mathbb{C}$ such that $E_{\{\lambda\}} = \mathrm{Id}$. Then $T = \lambda\mathrm{Id}$. $\square$

From this point on, every representation will be presumed to be a complex representation.

Corollary 1.3.8. *If G is an abelian topological group and $(\pi, \mathcal{H})$ is an irreducible unitary representation of G, then* $\dim \mathcal{H} = 1$.

Proof. Because G is abelian, we see that $\pi(g) \in \mathrm{U}(\mathcal{H})$ is an intertwining operator for fixed $g \in G$ (in particular, $\pi(g)\pi(h) = \pi(h)\pi(g)$ for all $h \in G$). It follows that $\pi(g) = \xi(g)\mathrm{Id}$ for some $\xi(g) \in \mathbb{C}$. Unitarity requires that $\xi(g) \in S^1 = \{z \in \mathbb{C} : |z| = 1\}$. Thus, we see that $\mathbb{C}v = \{\lambda v | \lambda \in \mathbb{C}\}$ is a closed invariant subspace of $\mathcal{H}$ for every $v \in \mathcal{H}$. Hence, the irreducibility of π requires that $\dim \mathcal{H} = 1$. $\square$

We remark that $\xi : G \to S^1$ from the above proof is a homomorphism. Such a homomorphism is called a **character** of G. Thus, the irreducible representations of G are, up to equivalence, simply the characters of G. We denote the set of all characters of an abelian group G by $\widehat{G}$. By using pointwise multiplication and pointwise convergence, it is possible to put an abelian topological group structure on $\widehat{G}$. While G and $\widehat{G}$ are not usually isomorphic to each other (the fact that $\mathbb{R}^n \cong \widehat{\mathbb{R}^n}$ is a rare exception), it is possible to construct a homomorphism

$$\begin{aligned} G &\to \widehat{\widehat{G}} \\ g &\mapsto \widehat{g} \end{aligned}$$

by setting $\widehat{g}(h) = h(g)$ for all $h \in \widehat{G}$. Pointryagin's famous duality theorem says that this homomorphism is in fact an isomorphism of topological groups if G is locally compact. While the theorem is not true in general for abelian groups which are not locally compact, it is still true for abelian groups which are constructed as direct limits of locally-compact abelian groups. But that is another story for another day. For now it suffices to say that this theory is basically "equivalent" to the Gelfand theory of commutative C^*-algebras.

We end this section by defining two very important classes of representations.

Definition 1.3.9. *A unitary representation* $(\pi, \mathcal{H})$ *of a topological group* G *is said to be* ***multiplicity free*** *if every decomposition* $\pi = \pi_1 \oplus \pi_2$ *of* π *into a direct sum of subrepresentations has the property that no subrepresentation of* π_1 *is equivalent to a subrepresentation of* π_2.

One can show that a unitary representation π is multiplicity-free if and only if its ring $\mathrm{Hom}(\pi, \pi)$ of intertwining operators is commutative. The term "multiplicity free" comes from the face that a direct sum $\pi = \oplus_{i \in I} \pi_i$ of irreducible representations of a group G is multiplicity free if and only if each equivalence class in $\widehat{G}$ appears at most once in the collection of π_i's. This basic result is a corollary of Schur's lemma (see [4, p. 123]).

Definition 1.3.10. *A unitary representation* $(\pi, \mathcal{H})$ *of a topological group* G *is said to be* ***primary*** *if the center of its ring of intertwining operators is trivial—that is, if*

$$Z(\mathrm{Hom}(\pi, \pi)) = \{\lambda \mathrm{Id} | \lambda \in \mathbb{C}\}.$$

One can show (see [4, p. 122]) that a direct sum $\pi = \oplus_{i \in I} \pi_i$ of irreducible representations of a group G is primary if and only if all the irreducible components π_i are equivalent to each other. However, for some groups it is possible to construct primary representations which cannot be decomposed into a direct sum of irreducible representations.

1.4 Fourier Analysis on $\mathbb{R}$: A Brief Reminder

In this section we briefly remind the reader of the basic results from Fourier analysis on $\mathbb{R}$. From our point of view, the "correct" way to view classical Fourier analysis is in terms of the decomopistion of a certain unitary representation of $\mathbb{R}$, so we begin by constructing some representations. Most of the material in this section is covered very well in the first half of [6].

First of all, Schur's lemma tells us that all irreducible representations have dimension one. In other words, we can identify the irreducible representations of $\mathbb{R}$ with its **unitary characters**, which are the continuous homomorphisms $\mathbb{R} \to S^1$, where $S^1 = \{z \in \mathbb{C} || z| = 1\}$. It is not difficult to see that the only unitary characters are those of the form $\theta_\xi : \mathbb{R} \to S^1$, where

$$\theta_\xi(x) = e^{2\pi i \xi x},$$

where $\xi \in \mathbb{R}$. Then θ_ξ may also be thought of as a unitary representation of $\mathbb{R}$ on the one-dimensional Hilbert space $\mathbb{C}$, where we define

$$\theta_\xi(x)z = e^{2\pi i \xi x} z$$

for all $z \in \mathbb{C}$.

Next, we define a representation $(L, L^2(\mathbb{R}))$, called the **regular representation** of $\mathbb{R}$, as follows. For each $y \in \mathbb{R}$, we set

$$(L(y)f)(x) = f(x-y).$$

for each $f \in L^2(\mathbb{R})$. In other words, L acts by shifting functions on $\mathbb{R}$ by translations. Because Lebesgue measure on $\mathbb{R}$ is translation-invariant, it is clear that $L(y)$ is a unitary operator on $\mathbb{L}^2(\mathbb{R})$ for every $y \in \mathbb{R}$:

$$||L(y)f||^2 = \int |f(x-y)|^2\, dx = \int |f(x)|^2\, dx = ||f||^2.$$

What is slightly less clear is that L is, in fact, a strongly continuous representation.

Lemma 1.4.1. *The representation $(L, L^2(\mathbb{R}))$ is strongly continuous.*

Proof. To show that L is strongly continuous, we fix $f \in L^2(\mathbb{R})$ and show that $\mathbb{R} \to L^2(\mathbb{R}), x \mapsto L(x)f$ is continuous. First, suppose that $x, y \in G$. Then

$$||L(x)f - L(y)f|| = ||L(y)(L(x-y)f - f)|| = ||L(x-y)f - f||,$$

where we use the fact that $L(y)$ is a unitary operator. From this it follows that it is sufficient to show that $x \mapsto L(x)f$ is continuous at $0 \in \mathbb{R}$.

Now fix $\epsilon > 0$. First we assume that $f \in C_c(\mathbb{R})$.Then there is $M > 0$ such that $\operatorname{supp} f \subseteq [-M, M]$. Next we recall that every compactly-supported function on a topological group is, in fact, *uniformly* continuous. That is, for each $\epsilon > 0$ there is a $\delta > 0$ of 0 such that $|f(x) - f(y)| < \epsilon/\sqrt{4M}$ for all $x, y \in \mathbb{R}$ such that $|x - y| < \delta$. In particular, we see that

$$||L(y)f - f||_\infty = \sup_{x \in \mathbb{R}} |f(x-y) - f(x)| < \epsilon/\sqrt{4M}$$

for all $y \in (-\delta, \delta)$. Without loss of generality, we assume that $\delta < M$. Then we have that $\operatorname{supp} L(y)f \subseteq [-2M, 2M]$.

$$||L(y)f - f||^2 = \int_{\mathbb{R}} |f(x-y) - f(x)|^2 dx \le \int_{[-2M,2M]} |\epsilon/\sqrt{4M}|^2 dx \le 4M \frac{\epsilon^2}{4M} = \epsilon^2.$$

and hence $||L(y)f - f|| < \epsilon$ for all $|y| < \delta$.

Extending the proof to general $f \in L^2(\mathbb{R})$ requires only using a standard $\epsilon/3$ argument and the fact that $C_c(\mathbb{R})$ is dense in $L^2(\mathbb{R})$. $\square$

1.4.1 Classical Fourier Analysis Results

We continue on to review the basic results of classical Fourier analysis. For $f \in L^1(G)$, we define the Fourier transform $\mathcal{F}(f) \equiv \widehat{f}$ of f to be a function on $\mathbb{R}$ defined by

$$\mathcal{F}(f)(\xi) \equiv \widehat{f}(\xi) = \int_{\mathbb{R}} f(x)e^{-2\pi i\xi x}dx \tag{1.1}$$

for $\xi \in \mathbb{R}$. One shows quickly that $\widehat{f} \in C_0(\mathbb{R})$ (that is, $\widehat{f}$ is continuous and decays to zero at infinity) for all $f \in L^1(\mathbb{R})$.

One of the basic tools for studying Fourier transforms is that of *convolution.* Recall that if $f, g \in C_c(\mathbb{R})$, then we define their **convolution** $f * g \in L^1(\mathbb{R})$ by

$$(f * g)(x) = \int_{\mathbb{R}} f(y)g(x-y)dy,$$

for all $x \in \mathbb{R}$ such that the integral is well-defined (one can show that the integral is well-defined for almost all $x \in \mathbb{R}$). Young's inequality shows that $||f * g||_1 \leq ||f||_1||g||_1$ for all $f, g \in L^1(\mathbb{R})$. One shows that the convolution product is bilinear and associative, and in fact turns $L^1(\mathbb{R})$ into a Banach algebra. In fact, $L^1(\mathbb{R})$ is a Banach-$*$ algebra if we use the involution $*$ on $L^1(\mathbb{R})$ defined by

$$f^*(x) = \overline{f(-x)}.$$

Unfortunately, $L^1(\mathbb{R})$ does not satisfy the equality $||f * f^*||_1 = ||f||^2$ for all $f \in L^1(\mathbb{R})$, and so it is not a C^*-algebra, but it is possible to use the norm on $L^1(\mathbb{R})$ to generate a $C - *$ algebra, called the *group C^*-algebra* of $\mathbb{R}$.

The Fourier transform satisfies a number of basic properties:

1. $\mathcal{F} : L^1(R) \to C_0(\mathbb{R})$ is linear
2. $\widehat{f * g}(\xi) = \widehat{f}(\xi)\widehat{g}(\xi)$ for all $f, g \in L_1(\mathbb{R})$
3. $\widehat{L(y)f}(x) = e^{2\pi i\xi y}\widehat{f}(\xi)$ for all $f \in L^1(\mathbb{R})$
4. $\widehat{\frac{d}{dx}f}(\xi) = 2\pi i\xi\widehat{f}(\xi)$ for all $f \in C_c^\infty(\mathbb{R})$
5. $\widehat{xf}(\xi) = -\frac{1}{2\pi i}\left(\frac{d}{d\xi}\widehat{f}\right)(\xi)$ for all $f \in C_c^\infty(\mathbb{R})$

Another important tool in Fourier analysis is the Schwartz space $\mathcal{S}(\mathbb{R})$ of rapidly decreasing smooth functions on $\mathbb{R}$, defined by

$$\mathcal{S}(\mathbb{R}) = \left\{ f \in C^\infty(\mathbb{R}) : \sup_{x \in \mathbb{R}} \left| x^p \frac{d^q}{dx^q} f(x) \right| < \infty \text{ for all } p, q \in \mathbb{N}_0 \right\}.$$

We give $\mathcal{S}(\mathbb{R})$ the topology of a Frechét space defined by the seminorms

$$||f||_{p,q} = \sup_{x \in \mathbb{R}} \left| x^p \frac{d^q}{dx^q} f(x) \right|$$

for all $p, q \in \mathbb{N}_0$. Using properties (4) and (5), it is possible to show that the Fourier transform extends to an isomorphism of vector spaces $\mathcal{F} : \mathcal{S}(\mathbb{R}) \to \mathcal{S}(\mathbb{R})$.

One next shows that if $f \in \mathcal{S}(\mathbb{R})$, then there is an inversion formula which recovers f from its Fourier transform $\widehat{f}$:

$$f(x) = \int_{\mathbb{R}} \widehat{f}(\xi) e^{2\pi i \xi x} d\xi. \tag{1.2}$$

for all $f \in \mathcal{S}(\mathbb{R})$. Furthermore, one shows that

$$\int_{\mathbb{R}} f(x)\widehat{g}(x)dx = \int_{\mathbb{R}} \widehat{f}(x)g(x)dx.$$

for all $f, g \in \mathcal{S}(\mathbb{R})$.

From this last result, it follows that the Fourier transform extends by continuity to a unitary operator $\mathcal{F} : L^2(\mathbb{R}) \to L^2(\mathbb{R})$. Furthermore, the inversion formula (1.2) holds for all $f \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. However, we remind the reader that the integral formula (1.1) holds only for $f \in L^1(\mathbb{R})$, because the integral may fail to converge for $f \in L^2(\mathbb{R})$.

1.4.2 Distributions and Plancherel Formula, Version I

Recall that a distribution on $\mathbb{R}$ is an linear functional on the space $\mathcal{D}(\mathbb{R}) \equiv \mathbb{C}_c^\infty(\mathbb{R})$. We denote the space of distributions by $\mathcal{D}'(\mathbb{R})$. (Recall that the space $\mathcal{D}(\mathbb{R})$ may be embedded *conjugate-linearly* as a dense subspace of $\mathcal{D}'(\mathbb{R})$). For a distribution $f \in \mathcal{D}'(\mathbb{R})$, we often use the bracket-notation

$$\langle \phi, f \rangle \equiv f(\phi)$$

for $\phi \in \mathcal{D}(\mathbb{R})$. We remind the reader that any locally integrable function $f : \mathbb{R} \to \mathbb{C}$ may be considered a distribution in $\mathcal{D}'(\mathbb{R})$ by setting

$$\langle \phi, f \rangle = \int_{\mathbb{R}} \phi(x) f(x) dx$$

for all $\phi \in \mathcal{D}(\mathbb{R})$. It is possible to take arbitrary derivatives of distributions, take convolutions with smooth functions, and even to take the convolution of two compactly-supported distributions.

For Fourier analysis, however, there is a special class of distributions which is especially useful, because the Fourier transform may be extended to this space. We denote the dual of $\mathcal{S}(\mathbb{R})$ by $\mathcal{S}'(\mathbb{R})$, which is said to be the space of **tempered distributions** on $\mathbb{R}$. Thus we have the following sequence of dense embeddings:

$$\mathcal{D}(\mathbb{R}) \subseteq \mathcal{S}(\mathbb{R}) \subseteq \mathbb{L}^2(\mathbb{R}) \subseteq \overline{\mathcal{S}'(\mathbb{R})} \subseteq \overline{\mathcal{D}'(\mathbb{R})},$$

where the overbars denote the fact that the embeddings of locally integrable functions into the distribution spaces are *conjugate linear* and not linear. While all locally integrable functions are distributions on $\mathbb{R}$, not all of them are *tempered* distributions. However, locally integrable functions of *slow growth* are tempered distributions. A function $f : \mathbb{R} \to \mathbb{C}$ is said to be of **slow growth** if

$$\sup_{x\in\mathbb{R}} |x+1|^{-p}|f(x)| < \infty$$

for some $p \in \mathbb{N}_0$. (The idea is that the functions grow no faster than a polynomial.) In particular, our characters θ_ξ, $\xi \in \mathbb{R}$ are tempered distributions.

We define the Fourier transform of a tempered distribution $f \in \mathcal{S}'(\mathbb{R})$ by setting

$$\left\langle \phi, \widehat{f} \right\rangle \equiv \left\langle \widehat{\phi}, f \right\rangle$$

for all $\phi \in \mathcal{S}(\mathbb{R})$. Furthermore, if the weak-$*$ topology is placed on $\mathcal{S}'(\mathbb{R})$, then the Fourier transform $\mathcal{F} : \mathcal{S}'(\mathbb{R}) \to \mathcal{S}'(\mathbb{R})$ becomes again an isomorphism of topological vector spaces.

There is another way to write down the Fourier inversion formula which will be useful to us in the future. Consider the Dirac delta function $\delta \in \mathcal{S}'(\mathbb{R})$ defined by

$$\langle \phi, \delta \rangle = \phi(0)$$

for all $\phi \in \mathcal{S}(\mathbb{R})$. But from the inversion formula, we know that

$$\begin{aligned}\langle \phi, \delta \rangle = \phi(0) = \int_{\mathbb{R}} \widehat{\phi}(\xi) d\xi &= \int_{\mathbb{R}}\int_{\mathbb{R}} \phi(x) e^{-2\pi i \xi x} dx d\xi \\ &= \int_{\mathbb{R}} \langle \phi, \theta_{-\xi} \rangle d\xi.\end{aligned}$$

In other words, it formally makes sense to write

$$\delta = \int_{\mathbb{R}} \theta_{-\xi} d\xi. \tag{1.3}$$

Furthermore, it is easy to recover the Fourier inversion formula from this decomposition of the delta distribution. If we define $\delta_y \in \mathcal{S}'(\mathbb{R})$ by $\langle \phi, \delta_y \rangle = \phi(y)$, then we see that

$$\langle \phi, \delta_y \rangle = \phi(y) = \int_{R} \langle \phi, \theta_{-\xi} \rangle \theta_x(\xi) d\xi$$

and hence we can formally write

$$\delta_x = \int_{\mathbb{R}} \theta_x(\xi)\theta_{-\xi} d\xi.$$

These decompositions of distributions as integrals of characters are one form of the *Plancherel Formula* for $\mathbb{R}$.

Next we show the other way of looking at the Fourier transform. Notice that there are two different representations which we can define on $L^2(\mathbb{R})$: one of them is the regular representation $(L, L^2(\mathbb{R}))$ which we have already defined. Another is the representation $(\pi, L^2(\mathbb{R}))$ given by

$$\pi(y)f(x) = \theta_x(y)f(x) = e^{2\pi i y x} f(x). \tag{1.4}$$

The Fourier transform $\mathcal{F} : L^2(\mathbb{R}) \to L^2(\mathbb{R})$ is, in fact, a unitary intertwining operator between this two representations since $\widehat{L(y)f}(x) = e^{2\pi i y x}\widehat{f}(x) = \pi(y)\widehat{f}(x)$.

1.4.3 Direct Integrals and Plancherel Formula, Version II

Now that we have access to the intertwining operator $\mathcal{F} \in \mathrm{Hom}_{\mathbb{R}}(L, \pi)$, we can describe the closed invariant subspaces of the regular representation $(L, \mathbb{L}^2(\mathbb{R}))$. In fact, if $\mathcal{K}$ is a closed L-invariant subspace of $\mathcal{K}$, then $\widehat{\mathcal{K}}$ is a closed π-invariant subspace of $L^2(\mathbb{R})$. One can then show that there is a Borel subset $A \subset \mathbb{R}$ such that $\mathcal{K} = L^2(\mathbb{R})_A$, where

$$L^2(\mathbb{R})_A = \{f \in L^2(\mathbb{R}) | \operatorname{supp} \widehat{f} \subseteq \overline{A}\},$$

where $\overline{A}$ denotes the closure of $A \subset \mathbb{R}$. Furthermore, it is easy to show that $L^2(\mathbb{R})_A \perp L^2(\mathbb{R})_B$ if $A \cap B$ has measure 0. However, if $L^2(\mathbb{R})_A \neq \{0\}$, then A is a set of nonzero measure. Thus there is a measurable subset $B \subset A$ such that both B and $A \backslash B$ have positive measure. Hecne,

$$L^2(\mathbb{R})_A = L^2(\mathbb{R})_B \oplus L^2(\mathbb{R})_{A \backslash B}$$

is a decomposition of $L^2(\mathbb{R})_A$ into a direct sum of nontrivial subrepresentations.

As we just saw, $(L, L^2(\mathbb{R}))$ does not have any nontrivial closed invariant subspaces, so it is impossible to write it as a direct sum of irreducible subrepresentations. nevertheless, it is possible to write L to be a "continuous orthogonal direct sum" (that is, a *direct integral*) of irreducible subrepresentations.

Fix a Hilbert space $\mathcal{H}$. Suppose that we have a measure space (X, μ). We say that a function $f : X \to \mathcal{H}$ is measurable if the the functions

$$X \to \mathcal{H}, x \mapsto \langle f(x), v \rangle$$

is measurable for all $v \in \mathcal{H}$. Then we define a new Hilbert space $\int_X^{\oplus} \mathcal{H} dx$ by setting

$$\int_X^{\oplus} \mathcal{H}\, dx = \left\{ f : X \to \mathcal{H} \,\middle|\, f \text{ is measurable and } \int_X ||f(x)||_{\mathcal{H}}^2 \, d\mu(x) < \infty \right\}$$

and then defining a new inner product by setting

$$\langle f, g \rangle = \int_X \langle f(x), g(x) \rangle_{\mathcal{H}} \, d\mu(x).$$

For example, it is easy to check that $\int_{\mathbb{R}}^{\oplus} \mathbb{C}\, dx = L^2(\mathbb{R})$.

Sometimes we will use the notation

$$f \equiv \int_X f(x) d\mu(x)$$

to emphasize the fact that we are considering f to be a sort of "continuous linear combination" of the values $f(x) \in \mathcal{H}$ for $x \in X$.

Next suppose that we have a topological group G and, for each $x \in X$, there is a representation $(\pi_x, \mathcal{H})$. We further suppose that the maps

$$x \mapsto \langle \pi_x(g)(f(x)), v \rangle$$

are measurable for all $v \in \mathcal{H}$. Then we define a new representation $\pi \equiv \int_G^{\oplus} \pi_x d\mu(x)$ on $\int_X^{\oplus} \mathcal{H} dx$ by setting

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

for all $g \in G$, $f \in \int_X^{\oplus} \mathcal{H} dx$, and $x \in X$. If we put this in our alternative notation, we have

$$\pi(g)\left(\int_X f(x)\, d\mu(x)\right) = \int_X \pi_x(g) f(x)\, d\mu(x).$$

Note that $\int_G^{\oplus} \pi_x d\mu(x)$ is a *unitary* representation of G on $\int_X^{\oplus} \mathcal{H} dx$ because

$$\begin{aligned} ||\pi(g)f||^2 &= \int_X ||\pi_x(g)f(x)||^2_{\mathcal{H}}\, d\mu(x) \\ &= \int_X ||f(x)||^2_{\mathcal{H}}\, d\mu(x) = ||f||^2, \end{aligned}$$

where we have used the fact that π_x is unitary for all $x \in X$. Strong continuity of $\int_G^{\oplus} \pi_x d\mu(x)$ follows from an elementary application of the Lebesgue Dominated Convergence Theorem.

We now see that the representation $(\pi, L^2(\mathbb{R}))$ defined in (1.4) is nothing more than the direct integral representation

$$\left(\int_{\mathbb{R}}^{\oplus} \theta_x\, dx, \int_R \mathbb{C}\, dx\right).$$

Finally, the equivalence of representations $(L, L^2(\mathbb{R})) \cong (\pi, L^2(\mathbb{R}))$ allows us to write

$$L \equiv \int_{\mathbb{R}}^{\oplus} \theta_x\, dx.$$

This decomposition of the regular representation into a direct integral of irreducible subrepresentations is also often referred to as a Plancherel formula.

This version of the Plancherel formula is often said to be "soft," however, because the equivalence class of a representation $\int_G^{\oplus} \pi_x d\mu(x)$ depends only on the measure class μ. To see this, suppose that σ is measure on X which is absolutely continuous with respect to μ. Then one can show that

$$\begin{aligned} T : \int_X^{\oplus} \mathcal{H}\, d\mu(x) &\to \int_X^{\oplus} \mathcal{H}\, d\sigma(x) \\ \int_X f(x) d\mu(x) &\mapsto \int_X f(x) \sqrt{\frac{d\mu(x)}{d\sigma(x)}}\, d\sigma(x) \end{aligned}$$

is a unitary intertwining operator in $\mathrm{Hom}_G\left(\int_X^{\oplus} \pi_x d\mu(x), \int_X^{\oplus} \pi_x d\sigma(x)\right)$.

Thus, we see that the equivalency

$$L \cong \int_{\mathbb{R}}^{\oplus} \theta_x\, d\mu(x).$$

also holds for any Borel measure μ on $\mathbb{R}$ which is absolutely continuous with respect to Lebesque measure. However, there is one property which sets apart the Lebesgue measure on $\mathbb{R}$: namely, Lebesque measure is the only measure on $\mathbb{R}$ such that the inversion formula is given by

$$f(x) = \int_X \widehat{f}(\xi)\theta_x(\xi)d\xi.$$

For this reason, we often refer to (1.3) as being the "strong" plancherel formula.

1.5 Positive-Definite Functions

In this section we explore the connection between unitary representations and positive-definite functions. We begin with a basic definition:

Definition 1.5.1. *Let G be a group. We say that a function $\phi : G \to \mathbb{C}$ is **positive-definite** if*

$$\sum_{i,j=1}^{n} \phi(g_i^{-1}g_j)c_i\overline{c_j} > 0$$

where $g_i \in G$ and $c_i \in \mathbb{C}$ for $1 \leq i \leq n$. In other words, we require that the matrix $[g_i^{-1}g_j]_{ij}$ be positive-definite for each choice of $g_1, \ldots, g_n \in G$ and $n \in \mathbb{N}$.

Positive-definite functions have several basic properties which may be proved directly from the definition (see [3, Lemma 5.1.8]):

Lemma 1.5.2. *If $\phi : G \to \mathbb{C}$ is a positive-definite function, then*

1. $\phi(e) > 0$
2. $|\phi(g)| \leq \phi(e)$ *for all* $g \in G$
3. $\phi(g^{-1}) = \overline{\phi(g)}$ *for all* $g \in G$

Proof. For example, to prove (1), we choose $n = 1$, $g_1 = e$, and $c_1 = 1$. It follows that $\phi(e) > 0$. To prove (3), we fix $g \in G$ and choose $n = 2$, $g_1 = g, g_2 = e$. Hence the matrix

$$\begin{pmatrix} \phi(e) & \phi(g) \\ \phi(g^{-1}) & \phi(e) \end{pmatrix}$$

is positive-definite. In particular, if we choose $c_1 = 1$ and $c_2 = i$, then we get the condition

$$2\phi(e) + i(\phi(g^{-1}) - \phi(g)) > 0$$

From this it follows that $\mathrm{Re}\,(\phi(g^{-1}) - \phi(g)) = 0$. Similarly, choosing $c_1 = c_2 = 1$, we obtain

$$2\phi(e) + \phi(g^{-1}) + \phi(g) > 0,$$

and hence $\mathrm{Im}\,(\phi(g^{-1}) + \phi(g)) = 0$. Thus, $\phi(g^{-1}) = \overline{\phi(g)}$. To prove (2), we note that since the matrix

$$\begin{pmatrix} \phi(e) & \phi(g) \\ \phi(g^{-1}) & \phi(e) \end{pmatrix} = \begin{pmatrix} \phi(e) & \phi(g) \\ \overline{\phi(g)} & \phi(e) \end{pmatrix}$$

is both self-adjoint and positive-definite, it must have a positive determinant. In other words, $\phi(e)^2 - |\phi(g)|^2 > 0$. □

The canonical examples of positive-definite functions are provided by matrix coefficients of unitary representations. That is, if $(\pi, \mathcal{H})$ is a unitary representation of a group G and $v \in \mathcal{H}\backslash\{0\}$, then the function $\phi_{\pi,v} : G \to \mathbb{C}$ given by

$$\phi_{\pi,v}(g) = \langle v, \pi(g)v \rangle \tag{1.1}$$

is continuous and positive-definite, as we now show straightforwardly using the unitarity of π and the definition of positive-definite functions. Continuity of $\phi_{\pi,v}$ follows immediately from the strong continuity of the representation π. The fact that it is positive-definite follows from the fact that if $g_i \in G$ and $c_i \in \mathbb{C}$ for $1 \le i \le n$, then

$$\begin{aligned} \sum_{i,j=1}^{n} \phi_{\pi,v}(g_i^{-1} g_j) c_i \overline{c_j} &= \sum_{i,j=1}^{n} c_i \overline{c_j} \langle \pi(g_i^{-1} g_j)v, v \rangle \\ &= \sum_{i,j=1}^{n} c_i \overline{c_j} \langle \pi(g_j)v, \pi(g_i)v \rangle \\ &= \left\langle \sum_{j=1}^{n} c_i \pi(g_j)v, \sum_{i=1}^{n} c_i \pi(g_i)v \right\rangle = \left\| \sum_{j=1}^{n} c_i \pi(g_j)v \right\|^2 > 0. \end{aligned}$$

Note that we used the unitarity of π in an essential way.

The key insight of Gelfand-Naimark-Segal is that every continuous positive-definite function arises in this way from a unitary representation. In particular, given a continuous positive-definite function $\phi : G \to \mathbb{C}$, one can define a representation. We now show how this may be done.

For each $g \in G$, define the function $g \cdot \phi : G \to \mathbb{C}$ by

$$g \cdot \phi(x) = \phi(g^{-1}x)$$

for each $x \in G$. We can then define the vector space

$$V_\phi = \langle \{ g \cdot \phi | g \in G \} \rangle,$$

which is the algebraic span of all G-translates of ϕ. We define a pre-Hilbert space structure on V_ϕ:

$$\left\langle \sum_{i=1}^{n} c_i (g_i \cdot \phi), \sum_{j=1}^{n} d_j (h_i \cdot \phi) \right\rangle = \sum_{i,j=1}^{n} \phi(g_i^{-1} h_j) c_i \overline{d_j} \tag{1.2}$$

where $c_i, d_j \in \mathbb{C}$ and $g_i, h_j \in G$. One needs to prove that this bilinear form is well-defined on V_ϕ (it is possible that an element of V_ϕ could be written as a linear combination of the translates of ϕ in more than way), but this can be done. It will then follow that V_ϕ is a pre-Hilbert space under $\langle , \rangle$.

We can then define a representation π_ϕ of G on V_ϕ by

$$\pi_\phi(g)v(h) = v(g^{-1}h)$$

for all $v \in V_f$ and $g, h \in G$. It is clear from (1.2) that π_ϕ extends to a unitary representation on the the Hilbert-space completion $\mathcal{H}_\phi$ of V_ϕ. Then one has

$$\phi(g) = \langle \phi, \pi(g)\phi \rangle_{\mathcal{H}_\phi}.$$

Thus every positive-definite function may be given the form (1.1). In fact, a stronger result may be proven:

Theorem 1.5.3. *(Gelfand-Naimark-Segal; see [3, p. 54, 61]). The map*

$$(\pi, v) \mapsto \phi_{\pi,v}$$

is a surjection from the set of all pairs (π, v) of cyclic representations $(\pi, \mathcal{H})$ of G and cyclic vectors $v \in \mathcal{H} \backslash \{0\}$ to the set of all continuous positive-definite functions on G.

Furthermore, suppose that $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{K})$ are unitary representations of G such that $v \in \mathcal{H}$ and $w \in \mathcal{K}$ are cyclic vectors. Then one has

$$\phi_{\pi,v} = \phi_{\sigma,w}$$

if and only if there is a unitary intertwining operator $T : \mathcal{H} \to \mathcal{K}$ such that $T(v) = w$.

Let G be a locally-compact topological group. We write $\mathcal{P}(G)$ for the space of all positive-definite functions ϕ on G such that $\phi(e) = 1$. One can show that $\mathcal{P}(G)$ is a closed convex subset of the space $L^\infty(G)$ of almost-everywhere-bounded measurable functions on G. The convexity may be shown by noticing that

$$\lambda\phi_{\pi,v} + (1-\lambda)\phi_{\sigma,w} = \phi_{\pi\oplus\sigma, \sqrt{\lambda}v + \sqrt{1-\lambda}w}, \tag{1.3}$$

where $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{K})$ are unitary representations of G with cyclic vectors $v \in \mathcal{H}$ and $w \in \mathcal{K}$.

In fact, $L^\infty(G)$ is the dual of the Banach space $L^1(G)$ by the Riesz Representation Theorem. One can show that $\mathcal{P}(G)$ is closed in the weak-$*$ topology on $L^\infty(G)$. Since $|\phi(g)| \leq \phi(e) = 1$ for all ϕ in $\mathcal{P}(G)$ and $g \in G$, we see that $\mathcal{P}(G)$ is contained in the unit ball $B_1(L^\infty(G))$. It follows from the Banach-Alaoglu theorem that $\mathcal{P}(G)$ is a compact convex subset of $L^\infty(G)$ in the weak-$*$ topology. Thus, the Krein-Milman theorem may be applied to $\mathcal{P}(G)$:

Theorem 1.5.4. *(Krein-Milman [3, Theorem 5.2.7]) If K is a compact, convex subset of a locally convex topological vector space V, then*

$$K = \overline{\mathrm{co}(\mathrm{ex}(K))},$$

where co *denotes the convex hull and* $\mathrm{ex}(K)$ *denotes the set of extremal points of* K*.*

In other words, all normalized positive-definite functions may be formed by taking a limit of convex combinations of normalized positive-definite functions. In fact, by exploiting the identity in (1.3), one has the following result:

Theorem 1.5.5. *Let G be a locally compact topological group. Then the extremal points of $\mathcal{P}(G)$ are given by functions of the form $\phi_{\pi,v}$, where $(\pi, \mathcal{H})$ is an irreducible representation of G and v is a cyclic unit vector in $\mathcal{H}$.*

Thus, positive-definite functions are generated in some sense by the ones coming from irreducible representations. These are just a few examples of how powerful theorems from functional analysis may be applied to provide insight into the decomposition of unitary representations.

1.6 Haar Measures

Though often conflated with each other, harmonic analysis and representation theory are not quite exactly the same fields of study. Representation theory studies, well, representations, while harmonic analysis studies analysis on groups and homogeneous spaces and is concerned specifically with understanding how to take general functions and break them down into easier-to-understand component parts. Nevertheless, harmonic analysis uses representation theory in an absolutely unavoidable way. For more details on the material in this section, I recommend Chapter 6 of [6] and Chapter 4 of [3].

First, we remind the reader that all locally-compact topological groups possess translation-invariant measures, called **Haar measures**. That is, for each locally-compact group G, there is a Radon measure μ_G such that

$$\int_G f(gx)d\mu_G(x) = \int_G f(x)d\mu_G(x) \tag{1.1}$$

for all $f \in C_c(G)$ and $g \in G$. Such measures, called **Haar measures**, are unique up to multiplication by a constant. If G is a compact group, then μ_G is a finite measure, which we will always normalize so that $\mu_G(G) = 1$.

A natural question to ask is how Haar measures behave under right-translation. Fix $g \in G$. Then we define a new measure μ_G^g on G by setting $\mu_G^g(A) = \mu_G(Ag)$ for each Borel subset $A \subseteq G$. Note that μ_G^g is also a left-invariant measure on G. Because left-invariant measures on G are unique up to constant multiple, there must be a real number $\Delta_G(g) > 0$ such that $\mu_G^g = \Delta_G(g)\mu_G$. Furthermore, one quickly checks that $\Delta_G(gh) = \Delta_G(g)\Delta_G(h)$ for each $g, h \in G$. We refer to the homomorphism $\Delta_G : G \to \mathbb{R}^+$ as the **modular function** for G. One can further show that Δ_G is continuous.

If $\Delta_G(g) = 1$ for all $g \in G$, then the Haar measure μ_G is both left-and right invariant, and we say that G is **unimodular**. It should be clear that the left-invariant Haar measure on an abelian group G must also be right-invariant. Furthermore, if G is compact, then the image of the modular function must be a compact subgroup of the multiplicative group $\mathbb{R}^+$. Since $\{1\}$ is the only compact subgroup of $\mathbb{R}^+$, it follows that $\Delta_G(g) = 1$ for all $g \in G$.

In addition to compact groups and abelian groups, it can be shown that all semisimple Lie groups and connected nilpotent Lie groups are unimodular [7, p. 88]. However, not all solvable groups are unimodular. While we are primarily concerned with studying semisimple Lie groups in this class, we will have to consider subgroups of semisimple groups which are *not* unimodular.

The existence of Haar measures has several important and useful consequences. For example, the next theorem shows that *every* representation of a compact group on a Hilbert space is equivalent to a unitary representation (this is not true for noncompact groups, as we will later see). One very important consequence of this result is that *every* finite-dimensional representation) of a compact group may be divided into a direct sum of irreducible representations.

Theorem 1.6.1. *(See also [17, Proposition 4.6]). If G is a compact topological group, then every norm-continuous representation $(\pi, \mathcal{H})$ of G on a Hilbert space is equivalent to a unitary representation.*

Proof. We denote the inner product on $\mathcal{H}$ by $\langle , \rangle_{\mathcal{H}}$ and construct a new inner product $\langle , \rangle_\pi$ on $\mathcal{H}$ by defining:

$$\langle v, w \rangle_\pi = \int_G \langle \pi(g)v, \pi(g)w \rangle_{\mathcal{H}} \, dg$$

for all $v, w \in \mathcal{H}$.

Now define

$$M = \sup_{g \in G} ||\pi(g)||_{\mathcal{H}}$$

and note that $M < \infty$ because π is norm-continuous and G is compact. We then have $||\pi(g)^{-1}||_{\mathcal{H}} < M$ for all $g \in G$. Thus

$$M^{-2}||v||^2_{\mathcal{H}} \leq ||v||^2_\pi = \int_G ||\pi(g)v||^2_{\mathcal{H}} \, dg \leq M^2||v||^2_{\mathcal{H}}$$

for all $v \in \mathcal{H}$. Hence the identity map on $\mathcal{H}$ forms a homeorphism between $\mathcal{H}$ under $\langle , \rangle_{\mathcal{H}}$ and $\mathcal{H}$ under $\langle , \rangle_\pi$.

Finally, for all $h \in G$ and $u, v \in \mathcal{H}$, we have that

$$\begin{aligned}\langle \pi(h)u, \pi(h)v \rangle_\pi &= \int_G \langle \pi(gh)v, \pi(gh)w \rangle_{\mathcal{H}} \, dg \\ &= \int_G \langle \pi(g)v, \pi(g)w \rangle_{\mathcal{H}} \, dg \\ &= \langle u, v \rangle_\pi .\end{aligned}$$

Thus, we see that π is a unitary representation of G on $\mathcal{H}$ under the inner product $\langle , \rangle_\pi$. □

The most important consequence of the existence of Haar measure, is that we can define a unitary representation of G which acts on a space of functions on G. It is this construction which provides the basic connection between harmonic analysis and representation theory. With a Haar measure dg on G, we may consider the Hilbert space $L^2(G)$. One can show that the action given by

$$(L(g)f)(x) = f(g^{-1}x) \tag{1.2}$$

for $g \in G$ and $f \in L^2(G)$ gives a continuous representation of G on $L^2(G)$ that is unitary by (1.1). This representation is called the **(left) regular representation** of G. We note that the proof of strong continuity is very similar to the proof for the regular representation of $\mathbb{R}$. The basic questions of Harmonic analysis, then, concern the decomposition of this representation into a direct sum or direct integral of irreducible subrepresentations.

We can also ask about G-invariant measures on homogeneous spaces: Suppose that H is a closed subgroup of G. Then G acts continuously and transitively on the left-coset space G/H, which we call a **homegeneous space**. It is then natural to ask whether there is a Borel measure $\mu_{G/H}$ on G/H such that

$$\int_{G/H} f(x \cdot gH) \, d\mu_{G/H}(gH) = \int_{G/H} f(gH) \, d\mu_{G/H}(gH)$$

for all $C_c(G/H)$ and $x \in G$. The answer is provided by the following theorem:

Theorem 1.6.2. *Suppose that H is a closed subgroup of a locally-compact topological group G. Then there exists a G-invariant measure on G/H if and only if $\Delta_G(h) = \Delta_H(h)$ for all $h \in H$ (in particular, this is always the case if H is a compact subgroup of G or if both G and H are unimodular). Furthermore, in that case the G-invariant measure is unique up to constant multiple and satisfies the property that*

$$\int_G f(g) \, dg = \int_{G/H} \int_H f(gh) \, dh \, d\mu_{G/H}(gH)$$

For the sake of clarity we also use the notation dx in place of $d\mu_{G/H}(x)$.

If H is a compact subgroup of G, then we can define the G-invariant measure on G/H by using the canonical quotient map $p : G \to G/H$ given by $p(g) = gH$. In fact, we see that $f \circ p \in C_c(G)$ for all $f \in C_c(G/H)$, and we can define the invariant measure dx on G/H by

$$\int_{G/H} f(x)dx = \int_G (f \circ p)(g)dg.$$

Once we have a G-invariant measure on $L^2(G/H)$, we construct a representation $(L, L^2(G/H))$ of G in basically the same way as the regular representation by setting

$$(L(g)f)(x) = f(g^{-1} \cdot x)$$

for all $g \in G$ and $x \in G/H$. Once again it is possible to show that this representation is strongly continuous. Unitarity follows immediately from the fact that the measure on G/H is invariant. This representation is called the **quasi-regular representation** of G on $L^2(G/H)$, and another basic problem of Harmonic analysis is the decomposition of this representation into irreducible components, when possible.

Even if there is no G-invariant measure on G/H, all is not lost, and it is still possible to define a quasi-regular representation. To begin, for each Borel measure $d\mu(x)$ on G/H and $g \in G$ we define a new measure $d\mu(g \cdot x)$ on G/H by setting

$$\int_{G/H} f(x)d\mu(g \cdot x) = \int_{G/H} f(g^{-1}x)d\mu(x)$$

for all $f \in C_c(G/H)$.

It is possible to show that there is always a unique class of Borel measures μ on G/H, called **quasi-invariant measures**, such that the the measures $d\mu(g \cdot x)$ and $d\mu(x)$ are absolutely continuous with respect to each other for all $g \in G$. In other words, for each $g \in G$ there is a Radon-Nikodym derivative given by

$$\rho(g, x) = \frac{d\mu(g \cdot x)}{d\mu(x)}$$

for each $g \in G$ and $x \in G/H$. In other words,

$$\int_{G/H} f(g^{-1} \cdot x)d\mu(x) = \int_G f(x)\rho(g, x)d\mu(x), \tag{1.3}$$

or, put in a different way,

$$\int_{G/H} f(g \cdot x)\rho(g, x)d\mu(x) = \int_G f(x)d\mu(x),$$

This function ρ satisfies two important properties, called the **cocycle identities**, which may be checked easily:

1. $\rho(e, x) = 1$ for almost all $x \in G/H$

2. $\rho(g_1 g_2, x) = \rho(g_2, x)\rho(g_1, g_2 \cdot x)$ for almost all $x \in G/H$ for each choice of $g_1, g_2 \in G$.

We can now construct the quasi-regular representation $(L, L^2(G/H))$ by setting

$$L(g)f(x) = \rho(g^{-1}, x)^{1/2} f(g^{-1} \cdot x).$$

We check that this gives a unitary representation by using the cocycle identities. First, unitarity follows because

$$\begin{aligned} ||L(g)f||^2 &= \int_{G/H} |\rho(g^{-1}, x)^{1/2} f(g^{-1} \cdot x)|^2 \, mu(x) \\ &= \int_{G/H} \rho(g^{-1}, x) |f(g^{-1} \cdot x)|^2 \, d\mu(x) \\ &= \int_{G/H} |f(x)|^2 \, d\mu(x) = ||f||^2 \end{aligned} \tag{1.4}$$

If $g_1, g_2 \in G$, then

$$\begin{aligned} L(g_1 g_2) f(x) &= \rho(g_2^{-1} g_1^{-1}, x)^{1/2} f(g_2^{-1} g_1^{-1} \cdot x) \\ &= \rho(g_1^{-1}, x)^{1/2} \rho(g_2^{-1}, g_1^{-1} \cdot x)^{1/2} f(g_2^{-1} g_1^{-1} \cdot x) \, d\mu(x) \\ &= L(g_1) \left(\rho(g_2^{-1}, \circ)^{1/2} f(g_2^{-1} \cdot \circ) \right) (x) \\ &= L(g_1) L(g_2) f(x) \end{aligned} \tag{1.5}$$

Finally, we leave it to the reader to show that any other choice of quasi-invariant measure μ' on G/H will produce a unitary representation equivalent to the one produced by μ. (This follows from the fact that there is a unique class of quasi-invariant measures on G/H, and thus μ' and μ must be absolutely continuous with respect to each other.)

1.7 General Definition of Direct Integrals

In Section 1.4.3, we introduced the concept of direct integrals of Hilbert spaces and representations over a measure space. In that section, however, we considered only the simpler situation in which all of the Hilbert spaces were the same. In this section, we take a Borel measure space (X, μ) and suppose that we have a Hilbert space $\mathcal{H}_x$ above each point $x \in X$. In other words, we have a fiber bundle over X where the fibers are all Hilbert spaces. The direct integral of the Hilbert spaces will then be the space of sections over this bundle. The tricky question is what sort of structure to put on this bundle. Should it be something like a smooth vector bundle or maybe merely a fiber bundle of topological spaces? In fact, we will essentially place only a Borel algebra on this bundle. This is the first example of a phenomenon in noncomutative harmonic analysis in which some objects should be thought of merely as Borel spaces and not necessarily as topological spaces or manifolds.

We now move to the details. Suppose that μ is a Borel measure on a topological space X, and that for each $x \in X$ we are given a unitary representation $(\pi_x, \mathcal{H}_x)$ of a group G. Suppose we are also given a collection of maps $s_i : X \to \dot{\cup}_{x\in X}\mathcal{H}_x$ for i in some countable index set I such that:

1. $s_i(x) \in \mathcal{H}_x$ for each $x \in X$ and $i \in I$.

2. $\langle s_i(x) | i \in I\rangle$ is dense in $\mathcal{H}_x$ for all $x \in X$.

3. $x \mapsto \langle s_i(x), s_j(x)\rangle_{\mathcal{H}_x}$ is a Borel-measurable function on X for all $i, j \in I$.

The set $\{s_i\}_{i\in I}$ is called a **measurable frame**. We then say that a map $s : X \to \dot{\cup}_{x\in X}\mathcal{H}_x$ is a **measurable section** if

1. $s(x) \in \mathcal{H}_x$ for each $x \in X$.

2. $x \mapsto \langle s(x), s_i(x)\rangle_{\mathcal{H}_x}$ is a Borel-measurable function on X for all $i \in I$.

Finally, we define a **direct-integral Hilbert space** by

$$\mathcal{H} \equiv \int_X^{\oplus} \mathcal{H}_x d\mu(x) = \left\{ \text{measurable sections } s \,\middle|\, \int_X ||s(x)||^2_{\mathcal{H}_x} d\mu(x) < \infty \right\}$$

where the inner product is given by

$$\langle u, v\rangle = \int_X \langle u(x), v(x)\rangle_{\mathcal{H}_x} d\mu(x)$$

for $u, v \in \mathcal{H}$. We can also define a continuous unitary representation $\pi \equiv \int_X^{\oplus} \pi_x d\mu(x)$ of G on $\mathcal{H}$ by

$$(\pi(g)s)(x) = \pi_x(g)(s(x))$$

for all $s \in \mathcal{H}$ and $g \in G$. We say that π is a **direct integral of the representations** $\mathcal{H}_x$ for $x \in X$.

1.8 Classes of Bounded Operators on a Hilbert Space

In this section, we introduce several different classes of bounded operators on a Hilbert space which will be of great use to use in the future. Some of the theorems are proved in detail, while others are only sketched or left to the reader. For more details, I highly recommend the book [6].

1.8.1 Compact Operators

Suppose that X and Y are Banach spaces. Then the space $\mathcal{L}(X,Y)$ of bounded linear operators from X to Y is also a Banach space under the operator norm. A bounded linear operator $T : X \to Y$ satisfies the property that $||Tx||_Y \leq ||T|| \cdot ||x||_X$ for all $x \in X$. In other words, changing x by a small amount only changes Tx by a small amount in the operator norm. In still other words, an operator $T : X \to Y$ is bounded if and only if the image under T of a bounded set in X is a bounded set in Y.

Sometimes, however, we want to impose an even stronger condition: that bounded sets in X get mapped to *precompact sets* in Y. This is in general a very strong condition, because compact sets in an infinite-dimensional vector space are "very small." In fact, if X is an infinite-dimensional Banach space, then its unit ball is not compact (you should prove this). In other words, even the identity operator $\mathrm{Id} : X \to X$ is not a compact operator if X is infinite-dimensional.

To simplify the discussion, we will assume that $\mathcal{H}$ is a Hilbert space and consider the space $\mathcal{B}(\mathcal{H})$ of bounded linear operators from $\mathcal{H}$ to $\mathcal{H}$, even though some of the basic theorems on compact operators generalize to the case of operators between two Banach spaces.

Definition 1.8.1. *Suppose that $T \in \mathcal{B}(\mathcal{H})$. We say that T is:* ***compact*** *if $T(U) \subseteq \mathcal{H}$ is precompact (i.e., has a compact closure) for every bounded set $U \subset \mathcal{H}$. We denote by $\mathcal{B}_c(\mathcal{H})$ the space of all compact operators from $\mathcal{H}$ to $\mathcal{H}$.*

Remark 1.8.2. Note that it if $T \in \mathcal{B}(\mathcal{H})$, then it is enough to show that $T(B_1(\mathcal{H})) \subset \mathcal{H}$. As another side remark, one must show that the compact operators form a linear subspace of $\mathcal{B}(\mathcal{H})$. We leave the proof of this statement to the reader.

Example 1.8.3. Suppose that $\mathcal{H}$ is a separable, infinite-dimensional Hilbert space.

1. We have already seen that $\mathrm{Id} : \mathcal{H} \to \mathcal{H}$ is not compact.

2. Suppose that $v \in \mathcal{H}$. Then the operator $P_v \in \mathcal{B}(\mathcal{H})$ defined by

$$P_v(w) = \langle w, v\rangle v$$

   for $w \in \mathcal{H}$ has a one-dimensional image. That is, the image of any vector in $\mathcal{H}$ under P_v will be a multiple of v. Since P_v is bounded, we see that bounded sets get mapped to bounded sets. But any bounded subset of $\mathbb{C}v$ must be compact. Thus P_v is compact.

3. More generally, suppose that $T \in \mathcal{B}(\mathcal{H})$ has a finite-dimensional image in $\mathcal{H}$. In this case, we say that T is a **finite-rank operator**. As in the previous example, since bounded subsets of a finite-dimensional subspace are always compact, it follows that T is in particular compact.

4. Next we present an example of an operator which has an infinite-dimensional image but is nevertheless compact. Consider ℓ^2, the Hilbert space of square-integrable sequences of complex numbers. This space has a canonical Hilbert space basis $\{e_i\}_{i\in\mathbb{N}}$ where e_i is the sequence which is 1 for the i^{th} index and zero for the others. We define an operator $T : \ell^2 \to \ell^2$ by $T(e_i) = (1/i)e_i$. That is,

$$T\left(\sum_{i=1}^{\infty} c_i e_i\right) = \sum_{i=1}^{\infty} \frac{1}{i} c_i e_i$$

for any $(c_1, c_2 \ldots) \in \ell^2$. One can check quickly that $||T|| = 1$ and so, in particular, $T \in \mathcal{B}(\mathcal{H})$.

To prove that T is, in fact, a compact operator, we must use the theorem (which we will prove shortly) which says that a limit (in the operator norm) of a sequence of compact operators is compact. To apply that theorem to this example, we define for each n a finite-rank operator $T_n : \mathcal{H} \to \mathcal{H}$ by $T_n = P_n T$, where P_n is the orthogonal projection of $\mathcal{H}$ onto the finite-dimensional subspace spanned by $\{e_1, \ldots, e_n\}$. In other words,

$$T_n(e_i) = \begin{cases} \frac{1}{i} e_i, & \text{if } i < n \\ 0, & \text{if } i > n. \end{cases}$$

We thus see that if $||x|| < 1$, then

$$\begin{aligned} ||(T - T_n)x|| &= \sum_{i=n}^{\infty} \left| \frac{1}{i} \langle x, e_i \rangle \right|^2 \\ &\leq \frac{1}{n} ||x||^2. \end{aligned}$$

Thus, it follows that $T_n \to T$ as $n \to \infty$ in the norm topology on $\mathcal{B}(\mathcal{H})$. Our next theorem then shows that T is compact. Notice that we could have replaced $\{1/i\}_{i\in\mathbb{N}}$ in this problem with any other sequence which approaches 0.

Theorem 1.8.4. *The space $\mathcal{B}_c(\mathcal{H})$ of compact operators on $\mathcal{H}$ is a closed subspace of $\mathcal{B}(\mathcal{H})$ under the operator norm.*

Proof. Suppose that $T_i \in \mathcal{B}_c(\mathcal{H})$ for each $i \in \mathbb{N}$ and that $\{T_i\}_{i\in\mathbb{N}}$ converges to a bounded operator $T \in \mathcal{B}(\mathcal{H})$ in the operator norm. We must show that T is compact.

Let $\{v_n\}_{n\in\mathbb{N}}$ be a sequence of vectors in the unit ball $B_1(\mathcal{H})$. Because T_1 is a compact operator, it follows that $T_1(B_1(\mathcal{H})) \subseteq \mathcal{H}$ is precompact, and thus there is a subsequence $\{v_{k_{1,n}}\}_{n\in\mathbb{N}}$ of $\{v_n\}_{n\in\mathbb{N}}$ such that $\{T_1(v_{k_{1,n}})\}_{n\in\mathbb{N}}$ is a Cauchy sequence. Similarly, because T_2 is a compact operator, we see that $T_2(B_1(\mathcal{H}))$ is precompact and so there is a subsequence $\{v_{k_{2,n}}\}_{n\in\mathbb{N}}$ of$\{v_{k_{1,n}}\}_{n\in\mathbb{N}}$ such that $\{T_2(v_{k_{2,n}})\}_{n\in\mathbb{N}}$ is Cauchy.

Repeating this process indefinitely, we obtain for each $m \in \mathbb{N}$ a subsequence $\{v_{k_{m+1,n}}\}_{n\in\mathbb{N}}$ of $\{v_{k_{m,n}}\}_{n\in\mathbb{N}}$ such that $\{T_i(v_{k_{i,n}})\}_{n\in\mathbb{N}}$ converges for all $i \leq m$. We then consider the "diagonal" subsequence $\{v_{k_{n,n}}\}_{n\in\mathbb{N}}$, which is a subsequence of all of the other subsequences. Thus, we see that $\{T_i(v_{k_{n,n}})\}_{n\in\mathbb{N}}$ is Cauchy $\mathcal{H}$ for all $i \in \mathbb{N}$.

Now fix $\epsilon > 0$. Because $T_i \to T$ in the norm topology, there is $N \in \mathbb{N}$ such that $||T - T_N|| < \epsilon$ for all $i > 0$. But there is also $M \in \mathbb{N}$ such that $||T_N(v_{k_{n,n}}) - T_N(v_{k_{m,m}})||$ for all $m, n \geq M$. We thus see that

$$\begin{aligned}||T(v_{k_{n,n}}) - T(v_{k_{m,m}})|| &\leq ||(T - T_N)(v_{k_{n,n}})|| + ||T_N(v_{k_{n,n}} - v_{k_{m,m}})|| \\ &+ ||(T - T_N)(v_{k_{m,m}})|| \\ &< 3\epsilon\end{aligned}$$

for all $m, n > M$. Thus, because $\mathcal{H}$ is complete, it follows that $\{T(v_{k_{n,n}}\}_{n\in\mathbb{N}}$ is a convergent subsequence of $\{T(v_n)\}_{n\in\mathbb{N}}$. Hence, $T(B_1(\mathcal{H}))$ is precompact, and thus T is compact. □

Remark 1.8.5. The previous theorem shows that we can readily produce compact operators by taking limits of finite-rank operators. Another well-known theorem shows that, in fact, *every* compact operator is a limit in the operator norm topology of a sequence of finite-rank operators (see [6]).

Next we collect some other properties of compact operators. One of the most important is that the compact operators form a two-sided ideal in $\mathcal{B}(\mathcal{H})$:

Theorem 1.8.6. *If $T \in \mathcal{B}_c(\mathcal{H})$ is a compact operator and $S \in \mathcal{B}(\mathcal{H})$ is a bounded operator, then TS and ST are compact operators.*

Proof. Student exercise. □

Remark 1.8.7. One consequence of the last two theorems is that, as a closed subalgebra of $\mathcal{B}(\mathcal{H})$ under the operator norm, $\mathcal{B}_c(\mathcal{H})$ is a C^*-algebra (that is, it is a Banach-$*$ algebra which satisfies the condition $||T^*T|| = ||T||^2$ for all $T \in \mathcal{B}_c(\mathcal{H})$). However, it is *not* a von Neumann algebra, because it is not a closed subalgebra in the weak topology.

Or, using the common definition of von Neumann algebra, we see that the centralizer of $\mathcal{B}_c(\mathcal{H})$ within $\mathcal{B}(\mathcal{H})$ is the trivial algebra $\mathbb{C}\mathrm{Id}$ of multiples of the identity matrix. Thus the double commutant of $\mathcal{B}_c(\mathcal{H})$ is all of $\mathcal{B}(\mathcal{H})$, and it follows that $\mathcal{B}_c(\mathcal{H})$ is not a von Neumann algebra. This underlines the intuitive idea that von Neumann algebras have to be "big enough" in some sense.

We do not have space to prove it here, but the following theorem is very important and is the key to proving the Peter-Weyl theorem for compact groups:

Theorem 1.8.8 (Spectral Theorem for Compact Self-Adjoint Operators)**.** *Suppose that $T \in \mathcal{B}_c(\mathcal{H})$ is a compact self-adjoint operator on $\mathcal{H}$. Then there are at most countably many eigenvalues in the spectrum of T, and the spectrum is* *either discrete or has a cluster point only at* 0*. Furthermore, each eigenspace* $\mathcal{H}_\lambda$ *for* $\lambda \in \operatorname{Spec}(T)$ *is finite-dimensional if* $\lambda \neq 0$ *and*

$$\mathcal{H} \cong \bigoplus_{\lambda \in \operatorname{Spec}(T)} \mathcal{H}_\lambda.$$

1.8.2 Hilbert-Schmidt Operators

In this section, we will talk about a class of operators which is even smaller than the class of compact operators. To motivate the definition, we note that if $\mathcal{H}$ is an infinite-dimensional separable Hilbert space with an orthonormal basis $\{e_1, e_2, \ldots\}$, then we can represent general bounded operators $T \in \mathcal{B}(\mathcal{H})$ by an "infinite matrix" $[A_{ij}]_{i,j \in \mathbb{N}}$, where the matrix components are given by $A_{ij} = \langle Ae_j, e_i \rangle$ for each $i, j \in \mathbb{N}$. In particular, we note that the j^{th} column in this matrix gives the coefficients of Ae_j with respect to the chosen basis, while the i^{th} row gives the complex conjugates of the coefficients of A^*e_i. Because the operator A is bounded, it follows in particular that

$$||Ae_j|| = \sum_{i=1}^{\infty} |A_{ij}|^2 < \infty$$

for each $j \in \mathbb{N}$. That is, each column in the infinite matrix forms a sequence in ℓ^2. Similarly, each row forms a sequence in ℓ^2.

In general, however, we cannot conclude that the matrix coefficients A_{ij}, taken altogether, form a square-summable sequence. For example, if we take the identity operator Id, then $\mathrm{Id}_{ij} = \delta_{ij}$, and hence $\sum_{i,j=1}^{\infty} |\mathrm{Id}_{ij}|^2 = \infty$. However, the class of bounded operators $A \in \mathcal{B}(\mathcal{H})$ such that $\sum_{i,j=1}^{\infty} |A_{ij}|^2 < \infty$ is very important, because we can put an inner product on this space which turns it into a Hilbert space.

Definition 1.8.9. *We say that a bounded operator* $A \in \mathcal{B}(\mathcal{H})$ *over a Hilbert Space* $\mathcal{H}$ *with orthonormal basis* $\{e_1, e_2, \ldots\}$ *is a* ***Hilbert-Schmidt operator*** *if*

$$||A||_{\text{HS}}^2 \equiv \sum_{i,j=1}^{\infty} |A_{ij}|^2 = \sum_{i,j=1}^{\infty} |\langle Ae_j, e_i \rangle|^2 = \sum_{j=1}^{\infty} ||Ae_j||^2 < \infty.$$

We denote the space of all Hilbert-Schmidt operators by $\mathcal{B}_{\text{HS}}(\mathcal{H})$*, and we say that* $||A||_{\text{HS}}$ *is the* ***Hilbert-Schmidt norm*** *of the operator* A*.*

The first thing which we have to do is show that this definition does not depend on the choice of basis. Suppose that $A \in \mathcal{B}_{\text{HS}}(\mathcal{H})$ and that$\{e_1, e_2, \ldots\}$

and $\{f_1, f_2, \ldots\}$ are orthonormal bases for $\mathcal{H}$. Then

$$\begin{aligned}\sum_{j=1}^{\infty} ||Ae_j||^2 &= \sum_{i,j=1}^{\infty} |\langle Ae_j, f_i\rangle|^2 \\ &= \sum_{i,j=1}^{\infty} |\langle e_j, A^* f_i\rangle|^2 \\ &= \sum_{i,j=1}^{\infty} |\langle A^* f_i, e_j\rangle|^2 \\ &= \sum_{i,j=1}^{\infty} |\langle A^* f_i, f_j\rangle|^2 = \sum_{i,j=1}^{\infty} |\langle A f_j, f_i\rangle|^2 = \sum_{j=1}^{\infty} ||Af_j||^2,\end{aligned}$$

where we have repeatedly used the fact that $||v||^2 = \sum_{i=1}^{\infty} |\langle v, e_i\rangle|^2 = \sum_{i=1}^{\infty} |\langle v, f_i\rangle|^2$ for all $v \in \mathcal{H}$. Thus, the definition does not depend on the choice of basis.

Remark 1.8.10. One must show that the Hilbert-Schmidt operators form a subspace of the bounded operators and that the Hilbert-Schmidt norm satisfies the triangle inequality. We leave this to the reader.

Theorem 1.8.11. *Every Hilbert-Schmidt operator is a compact operator.*

Proof. We sketch the proof here. Let $\mathcal{H}$ be a Hilbert space with orthonormal basis $\{e_1, e_2, \ldots\}$ as usual, and let $T \in \mathcal{B}_{\text{HS}}$ be a Hilbert-Schmidt operator. Consider for each n the orthogonal projection operator P_n which projects $\mathcal{H}$ onto the finite-dimensional subspace generated by $\{e_1, \ldots, e_n\}$. In particular, we see that the operator TP_n is a finite-rank operator for each $n \in \mathbb{N}$. We claim that $TP_n \to T$ as $n \to \mathbb{N}$ in the operator norm topology. It will then follow from Theorem 1.8.4 that T is a compact operator.

Fix $\epsilon > 0$. We notice that

$$||TP_n e_i||^2 = \begin{cases} ||Te_i||^2 & \text{if } i < n \\ 0 & \text{if } i \geq n \end{cases}$$

We also note that because $\sum_{i=1}^{\infty} ||Te_i||^2 < \infty$, the sequence $\{||Te_i||^2\}_{i\in\mathbb{N}}$ must converge to 0. Thus, there is $N \in \mathbb{N}$ such that $||Te_i||^2 < \epsilon$ for all $i \geq N$. Next suppose that $x \in \mathcal{H}$, and consider the orthogonal expansion $x = \sum_{j=1}^{\infty} c_j e_j$ where

$c_j \in \mathbb{C}$ for each $j \in \mathbb{N}$.

$$\begin{aligned}
||(T - TP_n)x||^2 &= \sum_{i=1}^{\infty} \left| \left\langle (T - TP_n) \sum_{j=1}^{\infty} c_j e_j, e_i \right\rangle \right|^2 \\
&= \sum_{i=1}^{\infty} \left| \left\langle \sum_{j=n}^{\infty} c_j T e_j, e_i \right\rangle \right|^2 \\
&= \sum_{j=n}^{\infty} \sum_{i=1}^{\infty} |c_j \langle T e_j, e_i \rangle|^2 \\
&= \sum_{j=n}^{\infty} |c_j|^2 ||T e_j||^2 \\
&\leq \epsilon \sum_{j=n}^{\infty} |c_j|^2 \leq \epsilon \sum_{j=1}^{\infty} |c_j|^2 = \epsilon ||x||^2.
\end{aligned}$$

Thus, $TP_n \to T$ in norm and we are done. □

Next we exam some important properties of Hilbert-Schmidt operators:

Theorem 1.8.12. *Suppose that $\mathcal{H}$ is a separable Hilbert space, that $S, T \in \mathcal{B}_{\mathrm{HS}}(\mathcal{H})$, and that $A, B \in \mathcal{B}(\mathcal{H})$. Then:*

1. $||T^*||_{\mathrm{HS}} = ||T||_{\mathrm{HS}}$.
2. $||ST||_{\mathrm{HS}} \leq ||S||_{\mathrm{HS}} \cdot ||T||_{\mathrm{HS}}$.
3. $||AT||_{\mathrm{HS}} \leq ||A|| \cdot ||T||_{\mathrm{HS}}$ *and* $||TA||_{\mathrm{HS}} \leq \cdot ||T||_{\mathrm{HS}} \cdot ||A||$, *where the norm on A is the usual operator norm.*

Remark 1.8.13. The first two properties in the above theorem show that $\mathcal{B}_{\mathrm{HS}}$ is a Banach-$*$ algebra under the norm $|| \cdot ||_{\mathrm{HS}}$, while the third shows that it is a two-sided ideal in $\mathcal{B}(\mathcal{H})$, although it is not a closed ideal because its closure under the norm topology turns out to be $\mathcal{B}_c(\mathcal{H})$.

Proof (of Theorem 1.8.12). We prove the second property and leave the others to the reader. We note that

$$\begin{aligned}
||ST||_{\mathrm{HS}}^2 &= \sum_{i,j=1}^{\infty} |\langle ST e_j, e_i \rangle|^2 \\
&= \sum_{i,j=1}^{\infty} |\langle T e_j, S^* e_i \rangle|^2 \\
&\leq \sum_{i,j=1}^{\infty} ||T e_j||^2 ||S^* e_i||^2 \\
&= \sum_{i=1}^{\infty} ||T e_j||^2 \sum_{j=1}^{\infty} ||S^* e_i||^2 \\
&= ||T||_{\mathrm{HS}}^2 ||S^*||_{\mathrm{HS}}^2 = ||S||_{\mathrm{HS}}^2 ||T||_{\mathrm{HS}}^2.
\end{aligned}$$

□

We next notice that one may put an inner product on $\mathcal{H}$ consistent with the norm $||\cdot||_{\mathrm{HS}}$. The basic idea is to take two operators $A, B \in \mathcal{B}_{\mathrm{HS}}(\mathcal{H})$ and consider the matrix coefficients $[A_{ij}]$ and $[B_{ij}]$ as sequences in ℓ^2, so that the inner product should be defined by

$$\langle A, B\rangle = \sum_{i,j=1}^{\infty} A_{ij}\overline{B_{ij}}.$$

We can simplify the above formula a bit:

$$\begin{aligned}\sum_{i,j=1}^{\infty} A_{ij}\overline{B_{ij}} &= \sum_{i,j=1}^{\infty}\langle Ae_j, e_i\rangle\langle e_i, Be_j\rangle \\ &= \sum_{i,j=1}^{\infty}\langle Ae_j, \langle Be_j, e_i\rangle e_i\rangle \\ &= \sum_{j=1}^{\infty}\langle Ae_j, Be_j\rangle.\end{aligned}$$

We leave it to the reader to prove that this definition indeed gives an inner product which turns $\mathcal{B}_{\mathrm{HS}}(\mathcal{H})$ into a Hilbert space. Because a given norm may be given by at most one inner product (you can use various polarization and parallelogram identities to recover the inner product from the norm), it follows that this inner product does not depend on the choice of basis.

There are several ways to look at Hilbert-Schmidt operators. Consider a Hilbert space $\mathcal{H}$. It turns out that there is a natural unitary isomorphism between the Hilbert space $\mathcal{B}_{\mathrm{HS}}(\mathcal{H})$ and the tensor product space $\mathcal{H}\otimes\mathcal{H}^*$, where $\mathcal{H}^*\equiv\overline{\mathcal{H}}$ is the linear dual of $\mathcal{H}$.

The definition of the tensor product requires some care. We can, of course, consider the algebraic tensor product $\mathcal{H}\otimes\mathcal{H}^*$, which consists of finite linear combinations of vectors of the form $v\otimes\overline{w}$, where $v, w\in\mathcal{H}$. The fact that we are considering the product $\mathcal{H}\otimes\mathcal{H}^*$ and not $\mathcal{H}\otimes\mathcal{H}$ means that the simple tensors will be linear in the first factor and conjugate-linear in the second factor (that is, $cv\otimes\overline{w} = c(v\otimes\overline{w})$ but $v\otimes c\overline{w} = \overline{c}(v\otimes w)$ for $c\in C$ and $v, w\in\mathcal{H}$). We then place an inner product on $\mathcal{H}\otimes\mathcal{H}^*$ such that

$$\langle v_1\otimes\overline{w_1}, v_2\otimes\overline{w_2}\rangle_{\mathcal{H}\otimes\mathcal{H}^*} = \langle v_1, v_2\rangle\langle w_2, w_1\rangle$$

for all pairs of simple tensors, and then we extend the definition by linearity. We remark to the reader that the ordering must be $\langle w_2, w_1\rangle$ and not $\langle w_1, w_2\rangle$ in the above definition in order for the inner product to be linear in the first variable and conjugate-linear in the second variable. In other words, the tensor product space $\mathcal{H}\otimes\mathcal{H}$ has an inner product given by the property

$$\langle v_1\otimes w_1, v_2\otimes w_2\rangle_{\mathcal{H}\otimes\mathcal{H}} = \langle v_1, v_2\rangle\langle w_1, w_2\rangle.$$

Of course, one must show that the inner product $\langle\, ,\, \rangle_{\mathcal{H}\otimes\mathcal{H}^*}$ defined above on simple tensors extends to a well-defined inner product on the space of all finite linear combinations of simple tensors. One then takes the Hilbert space-completion of this inner product to obtain the Hilbert space $\mathcal{H} \otimes \mathcal{H}^*$. As an example of a basis for this space, suppose that $\{e_1, e_2, \ldots\}$ and $\{f_1, f, 2, \ldots\}$ are both orthonormal bases for $\mathcal{H}$. Then $\{e_i \otimes \overline{f_j}\}_{i,j\in\mathbb{N}}$ forms an orthonormal basis for $\mathcal{H} \otimes \mathcal{H}^*$.

Finally, we demonstrate that there is a natural (i.e., coordinate-independent) unitary isomorphism between $U : \mathcal{H} \otimes \mathcal{H}^* \to \mathcal{B}_{HS}(\mathcal{H})$. First we consider a simple tensor $v \otimes \overline{w} \in \mathcal{H} \otimes \mathcal{H}^*$. We define $U(v \otimes \overline{w}) = P_{v,w}$ where $P_{v,w} \in \mathcal{B}(\mathcal{H})$ is the rank-one operator defined by $P_{v,w}(u) = \langle u, w\rangle v$ for all $u \in \mathcal{H}$. Next one shows that the expression $P_{v,w}$ is linear in v and conjugate-linear in w. One then shows, using the universal property of tensor products, that U extends to a well-defined linear isomorphism from the algebraic tensor product space $\mathcal{H}\otimes\mathcal{H}^*$ to the space of finite-rank operators. Finally, we note that the inner product agrees on rank-one operators and simple tensors:

$$\langle v_1 \otimes \overline{w_1}, v_2 \otimes \overline{w_2}\rangle_{\mathcal{H}\otimes\mathcal{H}^*} = \langle v_1, v_2\rangle\langle w_2, w_1\rangle = \langle P_{v_1,w_1}, P_{v_2,w_2}\rangle_{\mathcal{B}_{HS}(\mathcal{H})}.$$

An argument from linearity then shows that U extends to a unitary isomorphism from $\mathcal{H} \otimes \mathcal{H}^*$ to $\mathcal{B}_{HS}(\mathcal{H})$.

1.8.3 Trace-Class Operators

Recall that for a linear operator $L : \mathbb{C}^n \to \mathbb{C}^n$, the trace is defined to be the sum of the diagonal matrix coefficients: $\mathrm{Tr}(L) = \sum_{i=1}^n \langle Le_i, e_i\rangle$. Furthermore, the trace is linear, independent of the choice of basis, and satisfies the property that $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$ for all linear operators A, B on $\mathbb{C}^n$. It would be nice to extend this notion of trace to the infinite-dimensional context.

Suppose that $\mathcal{H}$ is a Hilbert space with an orthonormal basis $\{e_1, e_2, \ldots\}$. It is tempting to set $Tr(T) = \sum_{i=1}^n \langle Te_i, e_i\rangle$ for each bounded operator $T \in \mathcal{B}(\mathcal{H})$, but there is a problem—namely that this sum might diverge. Even worse, the sum might converge conditionally, so that the value depends not only on the choice of basis, but on the *ordering* of that basis! For instance, one can easily check that this series diverges for the identity operator (whose matrix is diagonal with ones along the diagonal) and converges conditionally for the operator whose matrix is diagonal and has diagonal entries given by the sequence $\{(-1)^n 1/n\}$.

In other words, we need to find a condition which guarantees that the sum $Tr(T) = \sum_{i=1}^n \langle Te_i, e_i\rangle$ converges absolutely. We begin with the following initial definition:

Definition 1.8.14. *If $A \in \mathcal{B}(\mathcal{H})$ is positive-definite, then we say that A is a **trace-class operator** if $\sum_{i=1}^n \langle Ae_i, e_i\rangle < \infty$. In that case, we say that the value of that sum is the **trace** of A and denote it by* $\mathrm{Tr}(A)$.

We remind the reader that a self-adjoint operator $A \in \mathcal{B}(\mathcal{H})$ is said to be **positive-definite** if $\langle Av, v\rangle > 0$ for all $v \in \mathcal{H}$. Thus, if A is positive-definite, then the series which defines the trace is a sum of positive numbers and thus either diverges to infinity or converges absolutely. Furthermore, we notice that if A is positive-definite, then

$$\mathrm{Tr}(A) = \sum_{i=1}^{n} \langle Ae_i, e_i \rangle = \sum_{i=1}^{n} \langle A^{1/2}e_i, A^{1/2}e_i \rangle = \langle A^{1/2}, A^{1/2} \rangle_{\mathrm{HS}} = ||A^{1/2}||_{\mathrm{HS}}^2,$$

Thus, we arrive at the very important conclusion that a positive-definite operator A is trace class if and only if its positive square root $A^{1/2}$ is Hilbert-Schmidt. We further conclude that the value of the trace is independent of the choice of basis.

More generally, suppose that $A \in \mathcal{H}(\mathcal{H})$. We define the **modulus** of A to be the operator $|A| = \sqrt{A^*A}$. (We remind the reader that the operator A^*A is always positive-definite and therefore has a unique positive-definite square root.) It is a well-known theorem that one can decompose the operator A as $A = |A|U$, where $|A|$ is positive-definite and U is a partial isometry (that is, U^*U and UU^* are orthogonal projections). This decomposition is known as the **polar decomposition**. In fact, in the case that $A \in \mathrm{GL}(\mathcal{H})$, then $|A| \in \mathrm{GL}(\mathcal{H})$ and one easily shows that $U = |A|^{-1}A$ is a unitary operator.

Definition 1.8.15. *If $A \in \mathcal{B}(\mathcal{H})$, then we say that it is a* ***trace-class operator*** *if the operator $|A|$ is trace-class in the sense of the previous definition and we set*

$$\mathrm{Tr}(A) = \sum_{i=1}^{\infty} \langle Ae_i, e_i \rangle$$

for the ***trace*** *of A.*

We need to show that the series defining $\mathrm{Tr}(A)$ is absolutely convergent for every trace-class operator and that its value is independent of the choice of basis. Take a trace-class operator A and write $A = |A|U$ where U is a partial isometry. In particular, $|A|^{1/2}$ is a Hilbert-Schmidt operator. Since U is a partial isometry, we see that $||U|| \leq 1$. Thus, $|A|^{1/2}U$ is also a Hilbert-Schmidt operator, and $|||A|^{1/2}U||_{\mathrm{HS}} \leq |||A|^{1/2}||_{\mathrm{HS}}$ by the basic properties of Hilbert-Schmidt operators. Then we see that

$$\begin{aligned}
|\mathrm{Tr}(A)| &= \left| \sum_{i=1}^{\infty} \langle |A|Ue_i, e_i \rangle \right| \\
&= \left| \sum_{i=1}^{\infty} \left\langle |A|^{1/2}Ue_i, |A|^{1/2}e_i \right\rangle \right| \\
&= \left| \left\langle |A|^{1/2}U, |A|^{1/2} \right\rangle_{\mathrm{HS}} \right| \\
&\leq ||\,|A|^{1/2}U\,||_{\mathrm{HS}} \cdot ||\,|A|^{1/2}\,||_{\mathrm{HS}} \\
&\leq ||\,|A|^{1/2}\,||_{\mathrm{HS}} \cdot ||\,|A|^{1/2}\,||_{\mathrm{HS}} = ||\,|A|^{1/2}\,||_{\mathrm{HS}}^2 \\
&= \mathrm{Tr}(|A|) < \infty
\end{aligned}$$

and in particular that

$$\mathrm{Tr}(A) = \left\langle |A|^{1/2}U, |A|^{1/2} \right\rangle_{\mathrm{HS}},$$

whose value does not depend on the choice of basis.

In the course of the previous proof, we notice that a trace-class operator may be written as a product of two Hilbert-Schmidt operators: namely, $A = |A|^{1/2} \cdot |A|^{1/2}U$. It thus follows that every trace-class operator is a Hilbert-Schmidt operator. More generally, if a bounded operator $T \in \mathcal{B}(\mathcal{H})$ can be written as $T = RS$, where both R and S are Hilbert-Schmidt operators, then one can show that T is a trace-class operator. To round out our observations of the connections with Hilbert-Schmidt operators, we note that a quick calculation shows that

$$\langle R, S \rangle_{\mathrm{HS}} = \mathrm{Tr}(S^* R).$$

for all Hilbert-Schmidt operators R and S.

Next, we note that it is possible to turn the set $\mathcal{B}_{\mathrm{Tr}}(\mathcal{H})$ of trace-class operators into a Banach-space. We leave the proof to the reader. First one must show that it is indeed a linear subspace of $\mathcal{B}(\mathcal{H})$, which is not too difficult. The norm on the trace-class operators is given by $||A||_{\mathrm{Tr}} = \mathrm{Tr}(|A|)$ for all $A \in \mathcal{B}_{\mathrm{Tr}}(\mathcal{H})$.

It can furthermore be shown that $\mathcal{B}_{\mathrm{Tr}}(\mathcal{H})$ is a two-sided ideal and that $||AB||_{\mathrm{Tr}} \leq ||A||_{\mathrm{Tr}} \cdot ||B||$ and $||BA||_{\mathrm{Tr}} \leq ||B|| \cdot ||A||_{\mathrm{Tr}}$ for all $A \in \mathcal{B}_{\mathrm{Tr}}(\mathcal{H})$ and $\mathcal{B}(\mathcal{H})$. Finally, one has the equality $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$. The reader is encouraged to prove these statements.

We have the following embeddings of two-sided ideals:

$$\mathcal{B}_{\mathrm{fin}}(\mathcal{H}) \subseteq \mathcal{B}_{\mathrm{Tr}}(\mathcal{H}) \subseteq \mathcal{B}_{\mathrm{HS}}(\mathcal{H}) \subseteq \mathcal{B}_c(\mathcal{H}) \subseteq \mathcal{B}(\mathcal{H})$$

where $\mathcal{B}_{\mathrm{fin}}(\mathcal{H})$ is the algebra of finite-rank operators. Each of these embeddings except the first and last are continuous dense embeddings of Banach spaces.

We round out this section by mentioning the so-called *Schatten p-classes* and corresponding norms. Fix $1 \leq p < \infty$. We say that an operator $A \in \mathcal{B}(\mathcal{H})$ is in the **Schatten p-class** if the operator $|A|^p$ is a trace-class operator (here we must use the functional calculus to define general positive powers of a positive-definite operator). We denote this class of operators by $\mathcal{B}_p(\mathcal{H})$ and give it the Banach-space norm

$$||A||_p = \mathrm{Tr}(|A|^p)^{1/p}$$

for all $A \in \mathcal{B}_p(\mathcal{H})$. The reader should quickly check that $\mathcal{B}_1(\mathcal{H}) = \mathcal{B}_{\mathrm{Tr}}(\mathcal{H})$ and $\mathcal{B}_2(\mathcal{H}) = \mathcal{B}_{\mathrm{HS}}(\mathcal{H})$. Furthermore, it is natural in some sense to set $\mathcal{B}_\infty \equiv \mathcal{B}_c(\infty)$ and set the norm $|| \cdot ||_\infty$ equal to the operator norm. One then has dense embeddings of Banach spaces $\mathcal{B}_p(\mathcal{H}) \subseteq \mathcal{B}_q(\mathcal{H})$ for all $p < q$. One can even prove a sort of Hölder inequality among these p-norms, which behave very similar to L^p spaces.

1.9 Integral Kernels for Hilbert-Schmidt Operators

In this section, we will show yet another way to view Hilbert-Schmidt operators in the context of function spaces. Suppose that (X, μ) is a σ-finite measure space and consider the Hilbert space $L^2(X, \mu)$. We say that an operator $T \in \mathcal{B}(\mathcal{H})$ is an **integral operator with kernel** K if $K : X \times X \to X$ is an almost-everywhere-defined function such that

$$(Tf)(x) = \int_X f(y)K(x,y)d\mu(y)$$

for all $f \in L^2(X)$. You should think of the kernel K as providing a "continuous collection of matrix coefficients." In general, the question of which kernel functions define bounded operators and which operators are given by integral kernels is a delicate one (see [6] for more), but here we will restrict ourselves to the case where $K \in L^2(X \times X)$. In fact, the main aim of this section is to motivate and sketch the proof of the following theorem:

Theorem 1.9.1. *There is a natural unitary isomorphism between* $L^2(X \times X)$ *and* $\mathcal{B}_{\mathrm{HS}}(L^2(X))$ *given by the correspondence* $K \mapsto T_K$ *for all* $K \in L^2(X \times X)$*, where* T_K *is the operator given by*

$$(T_K f)(x) = \int_X f(y)K(x,y)d\mu(y)$$

for all $f \in L^2(X)$*.*

Proof. The previous section provides a unitary map $L^2(X) \otimes L^2(X)^* \to \mathcal{B}_{\mathrm{HS}}(L^2(X))$. We will thus begin by constructing a unitary map $L^2(X \times X) \to L^2(X) \otimes L^2(X)^*$. The proof will be finished when we show that the composition of these two maps is equal to the map $K \mapsto T_K$ from $L^2(X \times X)$ to $\mathcal{B}_{\mathrm{HS}}(L^2(X))$.

Consider a simple tensor $f \otimes \overline{g} \in L^2(G) \otimes L^2(G)^*$, where $f, g \in L^2(G)$. We then map $f \otimes \overline{g}$ to the function in $L^2(X, Y)$ given by $(x, y) \mapsto f(x)\overline{g}(y)$. Because the value of this new function is linear in f and antilinear in g, we see that it extends to a linear map from $R : L^2(X) \otimes L^2(X)^* \to L^2(X \times X)$. The map is an isometry because:

$$\begin{aligned}\int_{X\times X} |f(x)\overline{g(y)}|^2 dx\, dy &= \int_X |f(x)|^2 dx \int_X |g(y)|^2 dy \\ &= ||f||^2||g||^2 \\ &= ||f \otimes \overline{g}||^2\end{aligned}$$

Thus, as soon as we show that R has a dense image, we will know that $R : L^2(X) \otimes L^2(X)^* \to L^2(X \times X)$ is a unitary map. Suppose that $\{e_i\}_{i\in\mathbb{N}}$ is an orthonormal basis for $L^2(X)$. Then so is $\{\overline{e_i}\}_{i\in\mathbb{N}}$. Thus $\{e_i \otimes \overline{e_j}\}_{i,j\in\mathbb{N}}$ is an orthonormal basis for $L^2(X) \otimes L^2(X)^*$. By an abuse of notation we will denote the function $(x,y) \mapsto e_i(x)\overline{e_j}(y)$ by $e_i \times \overline{e_j}$. To show that R has a dense image, it will be enough to show that $\{e_i \otimes \overline{e_j}\}_{i,j\in\mathbb{N}}$ is a basis for $L^2(X \otimes X)$. It is clear that this set of functions is orthonormal.

Suppose now that $f \in L^2(X \times X)$. For each $x, y \in X$, we use the notation f_y for the fuction given by $f_y(x) = f(x,y)$ and f^x for that given by $f^x(y) = f(x,y)$. Note that $f_y, f^x \in L^2(X)$ for almost all $x, y \in X$. We will also use the notation $\langle f_\bullet, e_i\rangle \in L^2(X)$ for the function $y \mapsto \langle f_y, e_i\rangle$. We will use Parseval's equality, which says that because $\{e_i\}_{i\in\mathbb{N}}$ is a basis for $L^2(X)$, one has that $||g||^2 = \sum_{i=1} \infty |\langle g, e_i\rangle|^2$ for all $g \in L^2(X)$. Then:

$$\begin{aligned}
\sum_{i,j=1}^{\infty} |\langle f, e_i \otimes \overline{e_j}\rangle|^2 &= \sum_{i,j=1}^{\infty} \left| \int_{X\times X} f(x,y)\overline{e_i(x)}e_j(y)dx\,dy \right|^2 \\
&= \sum_{i,j=1}^{\infty} \left| \int_X \langle f_y, e_i\rangle e_j(y)dy \right|^2 \\
&= \sum_{i,j=1}^{\infty} |\langle\langle f_\bullet, e_i\rangle, \overline{e_j}\rangle|^2 \\
&= \sum_{i=1}^{\infty} ||\langle f_\bullet, e_i\rangle||^2 \text{ (since } \{\overline{e_j}\}_{j\in\mathbb{N}} \text{ is a basis for } L^2(X)) \\
&= \sum_{i=1}^{\infty} \int_X |\langle f_y, e_i\rangle|^2 dy \\
&= \int_X ||f_y||^2 dy \text{ (since } \{e_i\}_{i\in\mathbb{N}} \text{ is a basis for } L^2(X)) \\
&= \int_X \int_X |f(x,y)|^2 dx\,dy = ||f||^2.
\end{aligned}$$

Hence $\{e_i \otimes \overline{e_j}\}_{i,j\in\mathbb{N}}$ is an orthonormal basis for $L^2(X \times X)$ by Pareseval's equality.

We thus have a unitary map $R : L^2(X) \otimes L^2(X)^* \to L^2(X \times X)$. Finally, we recall from the previous section that there is a unitary map $S : L^2(X) \otimes L^2(X)^* \to \mathcal{B}_{\mathrm{HS}}(L^2(X))$. Thus, the map $SR^{-1} : L^2(X \times X) \to \mathcal{B}_{\mathrm{HS}}(L^2(X))$ is the unitary map we want.

On simple tensors, one sees that $S(f \otimes \overline{g}) = P_{f,g}$ is the rank-one operator given by $P_{f,g}(h) = \langle h, g\rangle f$ for all $f, g \in L^2(X)$. Thus, if we consider the function $f \otimes \overline{g} \in L^2(X \times X)$ given by $(f \otimes \overline{g})(x,y) = f(x)\overline{g}(y)$, then we have that:

$$\begin{aligned}
SR^{-1}(f \otimes \overline{g})h(x) &= (P_{f,g}h)(x)dx \\
&= \langle h, g\rangle f(x)dx \\
&= \int_X h(y)f(x)\overline{g(y)}dy.
\end{aligned}$$

Hence, the simple tensor $f \otimes \overline{g} \in L^2(X \times X)$ corresponds to the integral kernel operator $T_{f\otimes\overline{g}}$ with kernel $f \otimes \overline{g}$.

Since the map $K \mapsto T_K$ is bilinear on simple tensors and agrees with SR^{-1} on all simple tensors, it follows that the map $K \mapsto T_K$ is exactly the same as the unitary operator SR^{-1}, and we are done. □

As an example of how operator kernels behave like matrices, we notice that if $K, R \in L^2(X \times X)$ are kernels for the Hilbert-Schmidt operators T_K and T_R, then $T_K T_R$ is a Hilbert-Schmidt operator with an integral kernel $S \in L^2(X \times X)$ which we now calculate:

$$\begin{aligned}(T_K T_R f)(x) &= T_K \left(\int_X f(y) R(\cdot, y) d\mu(y) \right)(x) \\ &= \int_X \int_X f(y) R(z, y) d\mu(y) K(x, z) d\mu(z) \\ &= \int_X f(y) \int_X K(x, z) R(z, y) d\mu(z) d\mu(y) \\ &= \int_X f(y) S(x, y) d\mu(y),\end{aligned}$$

where

$$S(x, y) = \int_X K(x, z) R(z, y) d\mu(z)$$

for almost all $x, y \in X$ (one must of course carefully justify the use of Fubini's Theorem). In other words, integral kernels multiply like matrices.

Next, we see how to compute the adjoint of an operator with an integral kernel. For each $K \in L^2(X, X)$, we define $K^* \in L^2(X, X)$ by setting $K^*(x, y) = \overline{K(y, x)}$. It then turns out that $T_K^* = T_{K^*}$. In fact, we have that:

$$\begin{aligned}\langle T_K f, g \rangle &= \int_X \int_X f(y) K(x, y) dy \, \overline{g(x)} dx \\ &= \int_X f(y) \int_X K(x, y) \overline{g(x)} dx \, dy \\ &= \int_X f(y) \overline{\int_X K^*(y, x) g(x) dx} \, dy \\ &= \langle f, T_{K^*} g \rangle\end{aligned}$$

for all $f, g \in L^2(X)$. (Of course, one must justify the use of Fubini's Theorem.)

In the future, we will also need the following theorem, which we will not prove.

Theorem 1.9.2. *Suppose that $T : L^2(X) \to L^2(X)$ is a Hilbert-Schmidt operator with kernel $K \in L^2(X \times X)$. If the map $x \mapsto K(x, x)$ is continuous almost everywhere and integrable (that is, an element of $L^1(X)$), then T is a trace-clase operator, and the trace is given by:*

$$\operatorname{Tr} T = \int_X K(x, x) dx.$$

This result is another example in which an integral kernel acts like a "continuous matrix."

1.10 Construction of New Representations from Old

In this section we introduce some of the operations which may be performed on representations to construct new ones. We have already seen two important examples: the direct sum of representations and its generalization, the direct integral.

1.10.1 The Contragredient of a Representation

Suppose that (π, V) is a continuous representation of a topological group G on a complex, locally convex topological vector space V. Let V^* denote the dual space of continuous linear functionals and give it a topology. If $\phi \in V^*$ and $v \in V$, then we will often use the bracket notation:

$$\langle v, \phi \rangle \equiv \phi(v)$$

for the evaluation of the functional ϕ. We wish to construct a "dual representation" to π, which should be a representation of G on V^*. To begin with, there is already a well-known process for taking operators in $\mathcal{B}(V)$ and dualizing them to produce operators from V^* to V^*. In particular, if $T : V \to V$ is a bounded operator, then we construct an operator $T^* : V^* \to V^*$ as follows. If $\phi \in V^*$, then $T^*\phi$ is the functional in V^* defined by

$$\langle v, T^*\phi \rangle \equiv \langle Tv, \phi \rangle. \tag{1.1}$$

In fact, one shows that $T^* : V^* \to V^*$ is continuous if V^* is given the weak-$*$ topology (and also if V^* is given the norm topology in the case that V is a Banach space).

We can now the **contragredient representation** (π^*, V^*) by setting

$$\pi^*(g) = \pi(g^{-1})^* \in \mathrm{GL}(V^*),$$

for each $g \in G$, where $\pi(g^{-1})^*$ denotes the adjoint of $\pi(g^{-1})$. Note that it is necessary to use g^{-1} in this definition instead of g in order to make $\pi^* : G \to \mathrm{GL}(V^*)$ a homomorphism, because both the operator adjoint and group inverses are anti-automorphism.

One thing that we didn't clarify above is what topology we wish to put on the dual space V^*. There are several choices, and it is natural to ask which topologies on V^* make the contragredient representation π^* continuous.

Theorem 1.10.1. *Suppose that (π, V) is a continuous representation of a topological group G on a complex, locally convex topological vector space V. Then the representation (π^*, V^*) is continuous if V^* is given the weak-$*$ topology.*

Proof. We begin by showing that that π^* is continuous if V^* is given the weak-$*$ topology. Fix $\lambda \in V^*$ and prove the continuity of the map $g \mapsto \pi^*(g)$. Because we have given V^* the weak-$*$ topology, it suffices to show that $G \to V$, $g \mapsto \langle v, \pi^*(g)\lambda\rangle$ is continuous for all $v \in V$. But $\langle v, \pi^*(g)\lambda\rangle = \langle \pi(g^{-1})v, \lambda\rangle$. The map $g \mapsto \pi(g^{-1})v$ is continuous by the strong continuity of π and because $g \mapsto g^{-1}$ is continuous. Finally, $\lambda : V \to C$ is continuous, and so it follows that $g \mapsto \langle \pi(g^{-1})v, \lambda\rangle = \langle v, \pi^*(g)\lambda\rangle$ is continuous for all choices of $v \in V$ and $\lambda \in V^*$. Thus π^* is a strongly continuous representation of G on V^*. □

Next we move on to the case of Banach spaces. Here, it is often natural to put a Banach space topology on V^* given by the operator norm, instead of the weak-$*$ topology. When V^* is given the norm topology, is π^* still a strongly continuous representation? The answer is, fortunately, yes.

Theorem 1.10.2. *Suppose that (π, V) is a continuous representation of a topological group G on a Banach space V. Then the representation (π^*, V^*) is continuous if V^* is given the norm topology.*

Proof. Left to the reader as an exercise. □

1.10.2 The Contragredient of a Unitary Representation

Finally, we move to the case of Hilbert spaces. Here we must be slightly careful, because there are two different notions of the adjoint of an operator floating around. In particular, if $\mathcal{H}$ is a complex Hilbert space and $T \in \mathcal{B}(\mathcal{H})$ is a bounded operator, then we can consider the operator T^* defined by (1.1), which is a bounded operator on $\mathcal{H}^*$. We can also consider the operator $T^* : \mathcal{H} \to \mathcal{H}$ defined so that for each $v \in \mathcal{H}$, T^*v is the vector in $\mathcal{H}$ such that

$$\langle w, T^*v\rangle_{\mathcal{H}} = \langle Tv, w\rangle_{\mathcal{H}} \tag{1.2}$$

for all $w \in \mathcal{H}$. In this case, however, the bracket is an inner product on $\mathcal{H}$ rather than the pairing between $\mathcal{H}$ and $\mathcal{H}^*$, as in (1.1).

If $\mathcal{H}$ were a Hilbert space over the field of *real* numbers, then there would be a natural linear isomorphism of $\mathcal{H}$ and $\mathcal{H}^*$. The problem is that for a *complex* Hilbert space, the natural identification of $\mathcal{H}$ with $\mathcal{H}^*$ is an *antilinear* isometry rather than a *linear* isometry. Of course, all separable infinite-dimensional Hilbert spaces are isomorphic, so there are, in fact, unitary isomorphisms between $\mathcal{H}$ and $\mathcal{H}^*$, but there is no canonical choice of such an isomorphism. There must be a choice of a *complex conjugation* on $\mathcal{H}$ before such an isomorphism is uniquely determined.

We briefly flesh out the details here. For each $v \in \mathcal{H}$, we associate $\overline{v} \in \mathcal{H}^*$ by setting $\overline{v}(w) = \langle w, v\rangle_{\mathcal{H}}$ for $w \in \mathcal{H}$. This map $\mathcal{H} \to \mathcal{H}^*$ given by $v \mapsto \overline{v}$ is sadly antilinear and not linear. In particular, we have that $\langle \overline{v}, \overline{w}\rangle_{\mathcal{H}^*} = \langle w, v\rangle_{\mathcal{H}}$.

Now suppose that we choose a real Hilbert space $\mathcal{K}$ such that $\mathcal{H} \cong \mathcal{K} \oplus i\mathcal{K}$. There are many such choices and in general there is no canonical choice, although we note that if $\mathcal{H}$ is a Hilbert space of functions with values in $\mathbb{C}$, then the real subspace of functions with values in $\mathbb{R}$ is a natural choice for $\mathcal{K}$. We say that $\mathcal{K}$ is a *real form* of $\mathcal{H}$. Under the choice of a real form $\mathcal{K}$, the *conjugation map* is the antiunitary operator $c : \mathcal{H} \to \mathcal{H}$ given by $c(v + iw) = v - iw$ if $v, w \in \mathcal{K}$. Then we may construct a unitary *linear* operator $U : \mathcal{H} \to \mathcal{H}^*$ by setting $v \mapsto \overline{c(v)}$ for each $v \in \mathcal{H}$ (it is linear because it is given by a composition of two antilinear operators).

Next suppose that $T : \mathcal{H} \to \mathcal{H}$. To avoid confusion, we donote the adjoint in the sense of (1.1) by $A^\dagger : \mathcal{H}^* \to \mathcal{H}^*$ and the adjoint in the sense of (1.2) by $A^* : \mathcal{H} \to \mathcal{H}$. We want to compare the operators $A^* : \mathcal{H} \to \mathcal{H}$ and $U^{-1}A^\dagger U : \mathcal{H} \to \mathcal{H}$, where $U : \mathcal{H} \to \mathcal{H}^*$ is the unitary identification of $\mathcal{H}$ with its dual. Now by definition, if $\phi \in \mathcal{H}^*$, then $A^\dagger\phi(v) = \phi(Av)$ for all $v \in \mathcal{H}$. In particular, if $w \in \mathcal{H}$, then $\overline{w} \in \mathcal{H}^*$ and thus $A^\dagger\overline{w}(v) = \overline{w}(Av) = \langle Av, w\rangle_{\mathcal{H}} = \langle v, A^*w\rangle$. In particular, we see that $A^\dagger(\overline{w}) = \overline{A^*(w)}$, where as usual the overline denotes the antilinear map from $\mathcal{H} \to \mathcal{H}^*$.

$$\begin{aligned}
\langle U^{-1}A^\dagger Uv, w\rangle_{\mathcal{H}} &= \langle A^\dagger Uv, Uw\rangle_{\mathcal{H}^*} \\
&= \langle A^\dagger\overline{c(v)}, \overline{c(w)}\rangle_{\mathcal{H}^*} \\
&= \langle \overline{A^*c(v)}, \overline{c(w)}\rangle_{\mathcal{H}^*} \\
&= \langle c(w), A^*c(v)\rangle_{\mathcal{H}} \\
&= \overline{\langle A^*c(v), c(w)\rangle_{\mathcal{H}}}
\end{aligned}$$

In other words, after using our unitary map $U_{\mathcal{H}} \to \mathcal{H}^*$ to pull back $A^\dagger : \mathcal{H}^* \to \mathcal{H}^*$ to a linear operator on $\mathcal{H}$, we see that the resulting operator is not quite the same as our old friend A^*. We will use the notation $U^{-1}A^\dagger U \equiv A^T$ for reasons which will soon become clear. In other words, the above calculations boil down to:

$$\langle A^T v, w\rangle_{\mathcal{H}} = \overline{\langle A^*c(v), c(w)\rangle_{\mathcal{H}}}.$$

Now suppose that $\{e_i\}$ is a basis for $\mathcal{H}$ which lies entirely inside $\mathcal{K}$. Then $c(e_i) = e_i$ for each basis element and thus we see that:

$$\langle A^T e_i, e_j\rangle_{\mathcal{H}} = \overline{\langle A^*e_i, e_j\rangle_{\mathcal{H}}}.$$

In other words, the matrix coefficients for A^T are the complex conjugates of the matrix coefficients for A^*. In fact, the matrix for A^T is simply the transpose of the matrix for A, where as the matrix for A^* is the conjugate transpose of that for A.

We are now finally able to discuss the contragredient of a unitary representation $(\pi, \mathcal{H})$ of G. Then there is the contragredient representation $(\pi^*, \mathcal{H}^*)$. We choose a real form $\mathcal{K}$ for $\mathcal{H}$ and use the unitary map $U : \mathcal{H} \to \mathcal{H}^*$ to pull π^* back to an equivalent representation on $\mathcal{H}$. Furthermore we pick a basis for $\{e_i\}$ is a basis for $\mathcal{H}$ which lies entirely inside $\mathcal{K}$. Recall that π^* is defined by $\pi^*(g) = \pi(g^{-1})^\dagger$. Thus, when we pull back to a representation on $\mathcal{H}$ (which we will denote by $\overline{\pi}$), we get that $\overline{\pi}(g) = \pi(g^{-1})^T$. Thus,

$$\langle\overline{\pi}(g)e_i, e_j\rangle = \overline{\langle\pi(g^{-1})^*e_i, e_j\rangle} = \overline{\langle\pi(g)e_i, e_j\rangle}$$

for each i, j and $g \in G$. In other words, $\overline{\pi}$ is the representation of G obtained from π by replacing all matrix coefficients with their conjugates. For this reason, we often describe $(\overline{\pi}, \mathcal{H})$ (and the equivalent representation $(\pi^*, \mathcal{H})$) by saying that it is the representation *conjugate* to π.

Example 1.10.3. Consider the one-dimensional representation χ_r of $\mathbb{R}$ given by $\theta_\xi(x)z = e^{2\pi i \xi x} z$ for all $z \in \mathbb{C}$, where $\xi \in \mathbb{R}$. Then the contragredient representation is given by taking complex conjugates of the matrix coefficients of θ_ξ. That is, $\overline{\theta_\xi}(x)z = e^{-2\pi i \xi x} z$ for all $z \in \mathbb{C}$, so that $\overline{\theta_\xi} = \theta_{-\xi}$. In particular, we see that θ_ξ is inequivalent to $\overline{\theta_\xi}$ in all cases except for the trivial representation, which satisfies $\overline{\theta_0} \equiv \theta_0$. We say that θ_0 is a **self-conjugate representation**.

Example 1.10.4. Now consider the group $\mathrm{SO}(n)$. Since $n \times n$ matrices with real coefficients may be naturally considered to be $n \times n$ matrices with complex coefficients, we see that $\mathrm{SO}(n)$ can be considered a subgroup of $\mathrm{U}(n)$. We can thus define a *natural representation* $(\pi, \mathbb{C}^n)$ of $\mathrm{SO}(n)$ by simply considering it as a subgroup of $\mathrm{U}(n)$—that is, it is a group of linear operators on $\mathbb{C}^n$, so we simply let those operators "act naturally" on $\mathbb{C}^n$. So π is the inclusion map $\pi : \mathrm{SO}(n) \to \mathrm{U}(n)$. However, we the matrix coefficients of any matrix in $\mathrm{SO}(n) \leq \mathrm{U}(n)$ are always real-valued. It follows that π is self-conjugate.

1.10.3 A Little More About Direct Sums

If (π, V) and (σ, W) are unitary representations, then we already know how to construct the direct sum $\pi \oplus \sigma$. In fact, if $\{(\pi_n, V_n)\}_{n \in \mathbb{N}}$ is a sequence of unitary representations of G, then the same definition may be modified in the obvious way to define the infinite direct-sum representation $(\oplus_{n=1}^\infty \pi_n, \oplus_{n=1}^\infty V_n)$. The main new result we wish to introduce in the section is the following:

Lemma 1.10.5. *Suppose that (π, V), (σ, W), and $(\rho, \mathcal{H})$ are unitary representations of a group G. Then there is a natural linear isomorphism:*

$$\mathrm{Hom}_G(\mathcal{H}, V \oplus W) \cong \mathrm{Hom}_G(\mathcal{H}, V) \oplus \mathrm{Hom}_G(\mathcal{H}, W).$$

In fact, the same equality holds true for infinite direct sums of representations.

Proof. Student exercise. □

1.10.4 Tensor Products of Representations: Inner and Outer

We have already defined the tensor product of two Hilbert spaces, and so we can now define the tensor product of unitary representations in the same way we defined the direct sum of unitary representations. Suppose that $(\pi, \mathcal{H})$ and $(\sigma, \mathcal{K})$ are unitary representations of a group G. Then we define a new representation $\pi \otimes \sigma$ of G on $\mathcal{H} \otimes \mathcal{K}$ by setting:

$$(\pi \otimes \sigma)(g)(v \otimes w) = \pi(g)v \otimes \sigma(g)w$$

for all $v \in \mathcal{H}$ and $w \in \mathcal{K}$. One must show that this representation is well-defined and extends by linearity from the simple tensors to all of $\mathcal{H} \otimes \mathcal{K}$ (in fact, this follows from the fact that on simple tensors $v \otimes w$, the value of $\pi(g)v \otimes \pi(g)w$ is bilinear in v and w). The fact that $\pi \otimes \sigma$ is unitary follows from the fact that on simple tensors, we have that:

$$||(\pi \otimes \sigma)(v \otimes w)||_{\mathcal{H}\otimes\mathcal{K}} = ||\pi(g)v||_{\mathcal{H}}||\sigma(g)w||_{\mathcal{K}} = ||v||_{\mathcal{H}}||w||_{\mathcal{K}}$$

for all $v \in \mathcal{H}$ and $w \in \mathcal{K}$. We say that $\pi \otimes \sigma$ is the **inner tensor product** or just **tensor product** of the representations π and σ.

This inner tensor product is a functor which takes two unitary representations of a group G and produces a new unitary representation of G on the tensor product of the Hilbert spaces. The *outer tensor product*, on the other hand, is a functor which takes a unitary representation of a group G and a unitary representation of another group H and then produces a new unitary representation of $G \times H$ on the tensor product of the Hilbert spaces.

In particular, suppose that $(\pi, \mathcal{H})$ is a unitary representation of G and that $(\sigma, \mathcal{K})$ is a unitary representation of H. Then we define a representation $\pi \boxtimes \sigma$ of the group $G \times H$ on the Hilbert space $\mathcal{H} \otimes \mathcal{K}$ by setting:

$$(\pi \boxtimes \sigma)(g, h)(v \otimes w) = \pi(g)v \otimes \sigma(h)w$$

for all $(g, h) \in G \times H$, $v \in \mathcal{H}$, and $w \in \mathcal{K}$. As before, one must show that this definition is well-defined and extends by linearity to all of $\mathcal{H} \otimes \mathcal{K}$. We say that $\pi \otimes \sigma$ is the **outer tensor product** of π and σ.

In general, the *inner* tensor product of two irreducible representations does of a group G does *not* produce an irreducible representation of G. For example, if we take a group G and a unitary representation (π, V) where $\dim V \geq 2$, then the exterior product $V \wedge V$ of anti symmetric tensors is an invariant subspace of $V \otimes V$. In fact, the orthogonal complement of $V \wedge V$ is the symmetric product $V \circledS V$. That is, one may write

$$V \otimes V \cong_G (V \wedge V) \oplus (V \circledS V).$$

We use the notation $(\pi \wedge \pi, V \wedge V)$ and $(\pi \circledS \pi, V \circledS V)$ to refer to the action of G on these invariant subspaces of $V \otimes V$. One may similar define an exterior power $\wedge^k \pi, \wedge^k V$ or symmetric power $\circledS^k \pi, \circledS^k V$ of a unitary representation (π, V).

However, *outer* tensor products behave very differently, as shown by the next theorem:

Theorem 1.10.6. *Suppose that $(\pi, \mathcal{H})$ is an irreducible unitary representation of a group G and that $(\sigma, \mathcal{K})$ is an irreducible unitary representation of a group H. Then $\pi \boxtimes \sigma$ is an irreducible representation of $G \times H$.*

Proof. By Schur's Lemma, it is sufficient to show that

$$\dim \operatorname{Hom}_{G\times H}(\pi \boxtimes \sigma, \pi \boxtimes \sigma) = 1. \tag{1.3}$$

In fact, by Schur's Lemma we already have that $\dim \mathrm{Hom}_G(\pi, \pi) = 1$ and that $\dim \mathrm{Hom}_H(\sigma, \sigma) = 1$.

Suppose that $T : \mathcal{H} \otimes \mathcal{K} \to \mathcal{H} \otimes \mathcal{K}$ is a bounded intertwining operator for $\pi \otimes \mathcal{K}$. Fix $w \in \mathcal{K}$ and consider the subspace $\mathcal{H} \otimes \{w\} = \{v \otimes w \in \mathcal{H} \otimes \mathcal{K} \mid v \in \mathcal{H}\}$. While $\mathcal{H} \otimes \{w\}$ is not in general an invariant subspace under the action of $G \otimes H$, it *is* invariant under the action of the subgroup $G \otimes \{e\}$. Furthermore, the restricted action of $G \otimes \{e\}$ on the subspace $\mathcal{H} \otimes \{w\}$ is equivalent to the representation $(\pi, \mathcal{H})$ and is in particular irreducible.

But T is an intertwining operator for the representation of $G \otimes H$, and it in particular intertwines the action of $G \cong G \otimes \{e\}$ on $\mathcal{H} \otimes \{w\}$. Since this subrepresentation is an irreducible representation of G, we see that there must be a constant $c(w) \in \mathbb{C}$ such that $T(v \otimes w) = c(w) v \otimes w$ for all $v \in \mathcal{H}$.

Performing the same argument for the restricted action of $\{e\} \otimes H$ on the subspace $\{v\} \otimes \mathcal{K} \subseteq \mathcal{H} \otimes \mathcal{K}$ for any fixed $v \in \mathcal{H}$, we see that there must be a constant $d(v) \in \mathbb{C}$ such that $T(v \otimes w) = d(v) v \otimes w$ for all $w \in \mathcal{K}$. Since we see that $T(v \otimes w) = d(v) v \otimes w = c(w) v \otimes w$ for all $v \in \mathcal{H}$ and $w \in \mathcal{K}$, it follows that we must have that $d(v) = c(w)$ for all $v \in \mathcal{H}, w \in \mathcal{K}$. In particular, there is a constant $c \in \mathbb{C}$ such that

$$T(v \otimes w) = cv \otimes w$$

for all $v \in \mathcal{H}$ and $w \in \mathcal{K}$. Extending by linearity to all of $\mathcal{H} \otimes \mathcal{K}$, we see that $T = c\,\mathrm{Id}$. Since T was arbitrary, we have thus shown (1.3). $\square$

Chapter 2

The Abstract Plancherel Formula: A Guiding Light

As we mentioned in the previous section, the foundational problem of harmonic analysis is to provide, for a particular group G or homogeneous space G/H, a decomposition of the regular representation into irreducible components. This is possible for a very broad class of locally-compact groups, called **Type I groups**.

Definition 2.0.7. *A topological group G is said to be of* ***Type I*** *if every primary representation of G decomposes into a direct sum of copies of the same irreducible representation. (See Definition 1.3.10 for the definition of primary representations.)*

This class includes all abelian groups, compact groups (see [7, p. 206]) and semisimple Lie groups (see [12, p. 230]), for example.

The Abstract Plancherel Theorem assures us that a decomposition of the regular representation is possible for all locally compact, separable Type I groups. To simplify the formulas, we will assume that G is also unimodular.

As with classical Fourier analysis, the convolution product is one of the most important tools in harmonic analysis on G. For any functions $f, g \in C_c(G)$, we define

$$f * g(x) = \int_G f(y) g(x^{-1}y)\, dy$$

Let $\widehat{G}$ denote the set of all equivalence classes of unitary representations of G. Then for each $f \in L^1(G)$ and each irreducible unitary representation $(\pi, \mathcal{H})$ we define the **operator-valued Fourier transform**

$$\widehat{f}(\pi) = \int_G f(g)\pi(g^{-1}) dg \in \mathcal{B}(\mathcal{H}).$$

This operator-valued integral is meant in the weak sense. That is,

$$\langle \widehat{f}(\pi)v, w\rangle = \int_G f(g)\langle \pi(g^{-1})v, w\rangle dg \in \mathcal{B}(\mathcal{H}).$$

for all $v, w \in \mathcal{H}$.

The operator-valued Fourier transform $\widehat{f}(\pi)$ is very closely related to the integrated representation of π, which is given by

$$\pi(f) = \int_G f(g)\pi(g)dg \in \mathcal{B}(\mathcal{H}).$$

In fact, it follows that $\widehat{f}(\pi) = \pi(f^\vee)$, where $f^\vee(g) = f(g^{-1})$ (here it is necessary to use that G is unimodular). Another quick computation shows that $\widehat{f}(\pi)^* = \pi(f)$. The Fourier transform and the integrated representations will both be useful in different contexts.

The operator-valued Fourier transform has all the expected properties of a Fourier transform:

1. $\widehat{f * g}(\pi) = \widehat{f}(\pi)\widehat{g}(\pi)$ and $\pi(f * g) = \pi(f)\pi(g)$ for $f, g \in L^1(G)$
2. $\widehat{f^*} = \widehat{f}(\pi)^*$ and $\pi(f^*) = \pi(f)^*$ for all $f \in L^1(G)$
3. $||\widehat{f}(\pi)||_{\text{op}} \leq ||f||_1$ and $||\pi(f)||_{\text{op}} \leq ||f||_1$, where $|| \cdot ||_{\text{op}}$ denotes the operator norm on $\mathcal{B}(\mathcal{H})$.

for $f, g \in L^1(G)$ and $\pi \in \widehat{G}$. Furthermore, if $g \in G$ and $f \in L^1(G)$, then

$$\begin{aligned}\widehat{L_g f}(\pi) &= \int_G f(g^{-1}x)\pi(x^{-1})dx \\ &= \int_G f(x)\pi((gx)^{-1})dx = \widehat{f}(\pi)\pi(g^{-1}).\end{aligned}$$

Similarly, on the level of integrated representations, we have that

$$\begin{aligned}\pi(L_g f) &= \int_G f(g^{-1}x)\pi(x)dx \\ &= \int_G f(x)\pi(gx)dx = \pi(g)\pi(f).\end{aligned}$$

We are now finally able to state the Abstract Plancherel Theorem. Its meaning will likely become clearer when we review the Peter-Weyl theory in the next section. It is most naturally stated as a decomposition of the biregular representation of $G \times G$ on $L^2(G)$ given by:

$$L(g_1, g_2)f(x) = f(g_1^{-1}xg_2)$$

for all $f \in L^2(G)$ and $x, g_1, g_2 \in G$.

Theorem 2.0.8. *(The Abstract Plancherel Theorem; see [4, p. 368]). Let G be a Type I separable, locally-compact unimodular topological group. For each $\pi \in \widehat{G}$,* *choose a representative irreducible representation* $(\pi, \mathcal{H}_\pi)$ *of* G*. Then there is a unique measure* μ *on* $\widehat{G}$ *such that*

$$L^2(G) \cong_{G\times G} \int_{\widehat{G}}^{\oplus} \mathcal{H}_\pi \otimes \overline{\mathcal{H}_\pi}\, d\mu(\pi)$$

and so that for all $f \in L^1(G) \cap L^2(G)$*, one has*

$$||f||^2 = \int_{\widehat{G}} \left|\left|\widehat{f}(\pi)\right|\right|_{HS}^2 d\mu(\pi),$$

where $||\cdot||_{HS}$ *denotes the Hilbert-Schmidt operator norm.*

Finally, it is possible to show that

$$f(g) = \int_{\widehat{G}} \operatorname{Tr}\left(\widehat{f}(\pi)\pi(g)\right) d\mu(\pi)$$

for all $g \in G$ *and all* f *may be written as a finite linear combination of convolutions of functions in* $L^1(G) \cap L^2(G)$*.*

Implicit in these statements, of course, is the claim that $\widehat{f}(\pi) \in \mathcal{B}(\mathcal{H})$ *is a Hilbert-Schmidt operator for* μ*-almost all* $\pi \in \widehat{G}$ *if* $f \in L^1(G) \cap L^2(G)$ *and that* $\widehat{f}(\pi)$ *is of trace class for* μ*-almost all* $\pi \in \widehat{G}$ *if* f *is a finite linear combination of convolutions of functions in* $L^1(G) \cap L^2(G)$*.*

To see the Plancherel Formula from a distributional point of view, now suppose that G is also a Lie group. Then we may consider the space $\mathcal{D}(G) \equiv C_c^\infty(G)$ of smooth, compactly supported functions on G and its dual space $\mathcal{D}'(G)$ of distributions on G. Then the delta distribution $\delta \in \mathcal{D}'(G)$ is defined, as usual, by

$$\langle \varphi, \delta \rangle = \varphi(e)$$

for all $\varphi \in \mathcal{D}(G)$. Next, hope to define a distribution $\theta_\pi \in \mathcal{D}'(G)$ for most representations $\pi \in \widehat{G}$ by setting

$$\langle \varphi, \theta_\pi \rangle = \operatorname{Tr}\left(\widehat{\varphi}(\pi)\right).$$

If, indeed, the operator $\widehat{\varphi} \in \mathcal{B}(\mathcal{H})$ is a trace-class operator for all $\varphi \in \mathcal{D}(G)$, then we say that π is a **representation with character** and say that θ_π is the **distributional character** of π.

Finally, the fact that $f(e) = \int_{\widehat{G}} \operatorname{Tr}\left(\widehat{f}(\pi)\right) d\mu(\pi)$ for f in a dense subspace of $L^2(G)$ leads us to expect that we might be able to decompose the delta distribution as

$$\delta = \int_{\widehat{G}} \theta_\pi \, d\mu(\pi).$$

This, in fact, will be the form in which we prove the Plancherel formula for semisimple groups.

2.1 The Plancherel Formula for Compact Groups: The Peter-Weyl Theory

In this section we discuss the Peter-Weyl theory of harmonic analysis on compact groups. The books [6, 7] are good sources for the material in this section. The theory is simpler for compact groups, because every Hilbert representation is unitarizable by Theorem 1.6.1 and furthermore, every unitary representation may be decomposed into a direct sum of finite-dimensional, irreducible subrepresentations:

Theorem 2.1.1. *Suppose that $(\pi, \mathcal{H})$ is a unitary representation of a compact group G. Then $\mathcal{H}$ contains a finite-dimensional invariant subrepresentation.*

Proof. Here we follow the proof in [7] For each $v \in \mathcal{H}$, we define an operator $K_v : \mathcal{H} \to \mathcal{H}$ by

$$K_v w = \int_G \langle w, \pi(g)v \rangle \pi(g) v dg.$$

Furthermore, one quickly checks that K is self-adjoint. The key to the proof is the fact that K is a compact operator on $\mathcal{H}$; we sketch the argument here. For each $u \in \mathcal{H}$, and $g \in G$ there is a self-adjoint, rank-one (and hence compact) operator P_u defined by $P_u w = \langle w, u \rangle u$. One then checks that

$$K_v = \int_G P_{\pi(g)v} dg$$

in the week sense. Because it is an integral of compact self-adjoint operators over a compact set, it follows that K_v is a compact, self-adjoint operator. Furthermore, one shows that $K_v \in \mathcal{B}(\mathcal{H})$ is an intertwining operator.

The spectral theorem for compact self-adjoint operators says that $\mathcal{H}$ decomposes into a discrete direct sum of eigenspaces of K_v, all of which are finite dimensional except possibly the 0 eigenspace. Each of these eigenspaces are invariant subspaces of $\mathcal{H}$ because K_v is an intertwining operator. If $v \neq 0$, then $K_v \neq 0$ and hence $\mathcal{H}$ possesses a finite-dimensional invariant subspace. □

Corollary 2.1.2. *Every irreducible representation of a compact group is finite-dimensional.*

We therefore restrict ourselves to finite-dimensional representations. Suppose that (π, V) is a finite-dimensional unitary representation of a compact group G. While it is no longer true that we can necessarily describe irreducible representations of G in terms of one-dimensional representations, there is still a useful notion of character. We define the **character** of π to be the function $\theta_\pi : G \to \mathbb{C}$ given by $\pi(g) = \mathrm{Tr}(\pi(g))$. Characters have the following properties, which may be verified without much difficulty. Suppose that (π, V) and (σ, W) are finite-dimensional representations of G. Then

1. $\theta_\pi(e) = d_\pi \equiv \dim V$

2. $\theta_{\pi\oplus\sigma} = \theta_\pi + \theta_\sigma$
3. $\theta_\pi = \theta_\sigma$ if and only if $\pi \cong \sigma$.

Similar in importance to characters are the **matrix coefficients** for a representation (π, V) of G. We set

$$\pi_{v,w}(g) = \langle v, \pi(g)w \rangle$$

for each $v, w \in V$. Note that if $\{e_1, \ldots e_{d_\pi}\}$ is a basis for V, then

$$\theta_\pi(g) = \mathrm{Tr}(\pi(g)) = \sum_{i=1}^{d_\pi} \pi_{e_i,e_i}(g).$$

Next, one proves the **Schur orthogonality relations**:

1. $\langle \pi_{v_1,w_1}, \pi_{v_2,w_2} \rangle = \frac{1}{d_\pi}\langle v_1, v_2\rangle\langle w_1, w_2\rangle$ for all $v_1, v_2, w_1, w_2 \in V$
2. $\langle \pi_{v_1,w_1}, \sigma_{v_2,w_2} \rangle = 0$ for all $v_1, w_1 \in V$ and $v_2, w_2 \in W$, if $\pi \ncong \sigma$

where (π, V) and (σ, W) are irreducible representations of G.

Suppose that (π, V_π) is an irreducible representation of G. Then the Schur orthogonality relations imply that $v \mapsto \pi_{v,w}$ is an injective intertwining operator from V to $L^2(G)$ for each fixed $w \in \mathcal{H}$. In fact, the space

$$\widetilde{\mathcal{H}}_\pi = \mathrm{span}\{\pi_{v,w} | v, w \in V\} \subseteq L^2(G)$$

may be naturally identified with the tensor product Hilbert space $V_\pi \otimes \overline{V_\pi}$, which of course may be naturally identified with the space $\mathcal{B}^2(V_\pi)$ of linear operators on V_π under the Hilbert-Schmidt inner product. The Peter-Weyl Theorem ties everything together:

Theorem 2.1.3. *(Peter-Weyl Theorem). Let G be a compact group, and let $\widehat{G}$ denote the set of all equivalence classes of irreducible unitary representations. For each $\pi \in \widehat{G}$, choose a representative irreducible representation (π, V_π) of G. Then*

$$L^2(G) \cong_{G\otimes G} \bigoplus_{\pi\in\widehat{G}} V_\pi \otimes \overline{V_\pi}$$

and so that for all $f \in L^2(G)$, one has

$$||f||^2 = \sum_{\pi\in\widehat{G}} d_\pi \left|\left|\widehat{f}(\pi)\right|\right|_{HS}^2,$$

where $||\cdot||_{HS}$ denotes the Hilbert-Schmidt operator norm.

Finally, it is possible to show that

$$f(g) = \sum_{\pi\in\widehat{G}} d_\pi \mathrm{Tr}\left(\widehat{f}(\pi)\pi(g)\right)$$

for all $f \in L^2(G)$, where the sum is taken in the L^2-norm.

The upshot of this theorem is that the Plancherel measure on $\widehat{G}$ is the discrete measure with weight d_π at each $\pi \in \widehat{G}$.

While the Plancherel measure for the groups we study will not be discrete, much of this theory does carry over to the so-called *discrete series* representations, as we will later see.

Chapter 3

Induction of Representations

Suppose that G and H are locally-compact topological groups and that H is a closed subgroup of G. We have already seen how to construct a "quasiregular" representation of G on $L^2(G/H)$, even in the case that G/H does not have invariant measures. In this section, we will see an important generalization of this representation, which will allow us to use a representation of H to *induce* a representation of G.

Before we present the definition, however, we remind the reader that there is a very simple way to take a representation of G and produce a representation of H. Suppose that $(\pi, \mathcal{H})$ is a unitary representation of G on a Hilbert space $\mathcal{H}$. Then we can construct a representation $(\operatorname{Res}_H^G \pi, \mathcal{H})$ of H by setting $\operatorname{Res}_H^G \pi(h) = \pi(h)$ for each $h \in H$. That is, $\operatorname{Res}_H^G \pi = \pi|_H$ is simply the restriction of the homomorphism $\pi : G \to U(\mathcal{H})$ to H. It is also not difficult to show that Res_H^G is a *functor* from the category of representations of G to the category of representations of H. (In particular, each intertwining operator $T : V \to W$, where (π, V) and (σ, W) are representations of G, is mapped to itself by the functor $\operatorname{Res}_H^G \pi$. This is because each such intertwining operator is automatically an intertwining operator between $\pi|_H$ and $\sigma|_H$.)

In this section, then, we will consider the dual concept: we will construct a functor from the category of representations of H to the category of representations of G. In fact, if H and G are compact groups, then we will see that this new functor is an *adjoint* of the restriction functor. This is the famous and ever-useful Frobenius Reciprocity Theorem. While this theorem is not true in the context of noncompact groups, there is a weaker replacement for it: the famous Imprimitivity Theorem of George Mackey, which forms the basis of the "little group method" or "Mackey machine."

We are now, finally, ready for the definition. We fix a unitary representation (σ, V) of H. Suppose that μ is a G-quasi-invariant measure on G/H. We will define a new representation $\operatorname{Ind}_H^G \sigma$ on a Hilbert space I^V which is defined as:

$$I^V = \left\{ f : G \to V \;\middle|\; \begin{array}{l} f(gh) = \sigma(h)^{-1} f(g) \text{ for all } g \in G, h \in H \\ \int_{G/H} ||f(g)||_V^2 d\mu(gH) < \infty \end{array} \right\},$$

in which the norm is given by

$$||f||^2 = \int_{G/H} ||f(g)||_V^2 d\mu(gH). \tag{3.1}$$

In other words, I^V is a space of functions on G with values in the Hilbert Space V. We remark to the reader that $||f(gh)||_V = ||\sigma(h)^{-1}f(g)||_V = ||f(g)||_V$ for all $g \in G$ and $h \in H$ because σ is a unitary representation on V. In other words, $g \mapsto ||f(g)||_V$ factors through to a map from G/H to $\mathbb{R}^+$. Thus the integral over G/H in (3.1) is well-defined.

To simplify the notation, for now we will set $\pi \equiv \operatorname{Ind}_G^H \sigma$. Since induction is intended to be a generalization of quasi-regular representations, it is perhaps not surprising that we set the action to be:

$$(\pi(g)f)(x) = \rho(g^{-1}, xH)^{1/2} f(g^{-1} \cdot x) \in V$$

for all $g \in G$ and $f \in I^V$ (recall that ρ is a function on $G \times G/H$ and that the norm of $f(x)$ depends only on the value of $xH \in G/H$).

To show that π is a representation, we should first check that $\pi(g)f$ is again an element of I^V for each $f \in I^V$. Suppose that $h \in H$, $g, x \in G$, and $f \in I^V$. Then

$$\begin{aligned}
(\pi(g)f)(xh) &= \rho(g^{-1}, xhH)^{1/2} f(g^{-1}xh) \\
&= \rho(g^{-1}, xH)^{1/2} \sigma(h)^{-1} f(g^{-1}h) \\
&= \sigma(h)^{-1}(\pi(g)f)(x),
\end{aligned}$$

by the equivariance property of $f \in I^V$, where we remind the reader that $\rho(g^{-1}, xH) \in \mathbb{R}^+$ and $f(g^{-1}xh) \in V$.

The fact that π gives a representation (i.e., that $\pi(g_1g_2) = \pi(g_1)\pi(g_2)$ for all $g_1, g_2 \in G$ follows from essentially the same calculation as in (1.5). The only difference is that we are dealing with V-valued functions on G instead of $\mathbb{C}$-valued functions on G/H. In fact, for each $g_1, g_2, x \in G$ and $f \in I^V$, we have:

$$\begin{aligned}
\pi(g_1g_2)f(x) &= \rho(g_2^{-1}g_1^{-1}, xH)^{1/2} f(g_2^{-1}g_1^{-1}x) \\
&= \rho(g_1^{-1}, xH)^{1/2} \rho(g_2^{-1}, g_1^{-1}xH)^{1/2} f(g_2^{-1}g_1^{-1}x)\, d\mu(x) \\
&= \pi(g_1) \left(\rho(g_2^{-1}, \circ H)^{1/2} f(g_2^{-1} \cdot \circ) \right)(x) \\
&= \pi(g_1)\pi(g_2)f(x).
\end{aligned}$$

That π is unitary follows from basically the same calculation in (1.4), except that we must use the fact that for each $f \in I^V$, the map $g \mapsto ||f(g)||_V$ factors through a function from G/H to $\mathbb{R}^+$. In fact, for each $g \in G$ and $f \in I^V$, we have

$$\begin{aligned}||\pi(g)f||^2 &= \int_{G/H} ||\rho(g^{-1},x)^{1/2} f(g^{-1}\cdot x)||_V^2 \, mu(x) \\ &= \int_{G/H} ||\rho(g^{-1},x)|f(g^{-1}\cdot x)|||_V^2 \, d\mu(x) \\ &= \int_{G/H} ||f(x)||_V^2 \, d\mu(x) = ||f||^2.\end{aligned}$$

We have now demonstrated that $\pi \equiv \mathrm{Ind}_H^G \sigma$ is a unitary representation of G.

In other words, we have a procedure for taking a representation of the smaller group H and constructing a representation of the bigger group, G. In fact, this procedure is a *functor* from the category of representations of G to the category of representations of H. There is some work to do, however, to prove this claim. First of all, a functor should not only take objects in one category to objects in the other, but also morphisms from the one category to morphisms for the other category. In other words, for each pair (σ, V) and (π, W) of unitary representations of H and each intertwining operator $T : V \to W$, we must produce in a natural way an intertwining operator $\mathrm{Ind}_H^G T : I^V \to I^W$. In fact, this is not so difficult. Supposing that $T \in \mathrm{Hom}_H(V, W)$ and $f \in I^V$, then $(\mathrm{Ind}_H^G T)f$ should be in I^W (in particular, it should be a W-valued function on G). We define $(\mathrm{Ind}_H^G T)f$ by setting

$$(\mathrm{Ind}_H^G T)f(x) \equiv T(f(x)) \in W.$$

for each $x \in G$.

To check that Ind_H^G is a functor, we need to show that it sends identity morphisms to identity morphisms and compositions of morphisms to compositions of morphisms. First of all, if (σ, V) is a unitary representation of H and $1_V : V \to V$ is the identity map, then $\mathrm{Ind}_H^G 1_V : I^V \to I^V$ is given by

$$(\mathrm{Ind}_H^G 1_V)f(x) = 1_V(f(x)) = f(x)$$

for each $f \in I^V$ and $x \in G$. Thus, we see that $(\mathrm{Ind}_H^G 1_V) = 1_{I^V}$.

Next, suppose that (σ_1, V_1), (σ_2, W), and (σ_3, V_3) are unitary representations of H and that $T : V_2 \to V_3$ and $S : V_1 \to V_2$ are H-intertwining operators. We must show that $\mathrm{Ind}_H^G(T \circ S) = (\mathrm{Ind}_H^G T) \circ (\mathrm{Ind}_H^G S)$. To do that, we fix $f \in I^{V_1}$ and $x \in G$. Then

$$(\mathrm{Ind}_H^G TS)f(x) \equiv T(S(f(x))) = \mathrm{Ind}_H^G T(\mathrm{Ind}_H^G S(f))(x),$$

and we are done. Thus Ind_H^G is a functor, which means that our method of producing representations of G from representations of H is a good one, from the perspective of category theory.

Our next steps will be to examine the simpler case in which both G and H are compact groups. In that case, we will be able to show that Res_H^G and Ind_H^G

are *adjoint functors*. This is, of course, the famous *Frobenious Reciprocity*. Next we will show that the induction functor commutes with direct sums and that so-called *induction in stages* works. We will then show two other routes to the construction of induced representations. We will end with some notes about the construction of intertwining operators between induced representations for finite groups, which is the inspiration for the theory of intertwining operators between generalized principal series representations.

3.1 Frobenius Reciprocity

In this section, we suppose that G and H are compact groups. Then G/H has an invariant-measure and, in fact, the definition of induction simplifies in the following way: If (σ, V) is an irreducible unitary representation of H, then the induced representation $(\operatorname{Ind}_H^G \sigma, I^V)$ is given by:

$$I^V = \left\{ f : G \to V \;\middle|\; \begin{array}{l} f(gh) = \sigma(h)^{-1} f(g) \text{ for all } g \in G, h \in H \\ \int_G ||f(g)||_V^2 dg < \infty \end{array} \right\},$$
$$(\operatorname{Ind}_H^G \sigma)(g) f(x) = f(g^{-1}x) \in V \text{ for } f \in I^V \text{ and } g, x \in G.$$

Now because G is compact, it follows that every continuous function from G to V is square-integrable, and in particular I^V admits the dense invariant subspace:

$$I_c^V = \left\{ f : G \to V \;\middle|\; \begin{array}{l} f(gh) = \sigma(h)^{-1} f(g) \text{ for all } g \in G, h \in H \\ f \text{ continuous} \end{array} \right\}.$$

The advantage of this dense subspace is that, unlike square-integrable functions (which are actually equivalence classes of functions) it is actually possible to evaluate a continuous function at a point. With this dense subspace it is possible to prove the Frobenius Reciprocity theorem:

Theorem 3.1.1. *Let G be a compact group with a closed subgroup H, and suppose that and (σ, V) and (π, W) are finite-dimensional unitary representations of H and G, respectively. Then there is an isomorphism of vector spaces*

$$\operatorname{Hom}_H(\operatorname{Res}_H^G \pi, \sigma) \cong \operatorname{Hom}_G(\pi, \operatorname{Ind}_H^G \sigma).$$

In other words, Ind_H^G and Res_H^G are adjoint functors. One can, of course, write the same statement in terms of vector spaces:

$$\operatorname{Hom}_H(W, V) \cong \operatorname{Hom}_G(W, I^V).$$

Proof. We begin by defining a linear map

$$R : \operatorname{Hom}_H(W, V) \to \operatorname{Hom}_G(W, I^V)$$
$$T \mapsto R_T.$$

The proof of this theorem is a game of playing with the various equivariance properties of intertwining operators and functions in induced representations.

Fix $T \in \mathrm{Hom}_H(W, V)$. In particular, for each w, we set $R_T w$ to be the function from G to V:

$$(R_T w)(g) = T(\pi(g^{-1})w) \in V.$$

It is clear that this function is continuous (i.e., $R_T w \in I_c^V$) by the strong continuity of π. It follows that $R_T w \in L^2(G)$. Note that we are here using the compactness of G and H. That it satisfies the equivariance property for membership in I^V follows from:

$$(R_T w)(gh) = T(\pi(h^{-1}g^{-1})w) = \sigma(h^{-1})T(\pi(g^{-1})w) = \sigma(h^{-1})R_T w(g),$$

where $g \in G$ and $h \in H$, since T intertwines π and σ as representations of H.

In other words, $R_T w \in I^V$ for all $w \in W$. It remains to be shown that R_T intertwines π and $\mathrm{Ind}_H^G \sigma$. In fact,

$$\begin{aligned}
R_T(\pi(g)w)(x) &= T(\pi(x^{-1})\pi(g)w)) = T(\pi(g^{-1}x)^{-1}w) \\
&= (R_T w)(g^{-1}x) \\
&= [(\mathrm{Ind}_H^G)\sigma(g)(R_T w)](x)
\end{aligned}$$

for all $w \in W$ and $g, x \in G$. Thus R_T is an intertwining operator.

Finally, we show that R_T is a bounded operator. Suppose that $w \in W$. Then $||R_T w(g)||_V = ||T(\pi(g)w)||_V \le ||T||||\pi(g)w||_V = ||T||||w||_V$ for all $g \in G$ and $w \in W$. It follows that

$$||R_T w||^2 = \int_G ||T(\pi(g^{-1})w)||_V^2 dg \le \int_G ||T||^2||w||_V^2 dg = ||T||^2||w||_V^2$$

for each $w \in W$, so that $||R_T|| \le ||T||$. Here we have once again used that G is compact. Thus R_T is a bounded operator and, in particular, $R_T \in \mathrm{Hom}_G(W, I^V)$.

It is clear that the map $R : T \mapsto R_T$ is linear. It remains only to construct a linear inverse to R:

$$\begin{aligned}
S : \mathrm{Hom}_G(W, I^V) &\to \mathrm{Hom}_H(W, V) \\
T &\mapsto S_T.
\end{aligned}$$

Suppose that $T \in \mathrm{Hom}_G(W, I^V)$. We would like to define S_T by setting $S_T w = (T(w))(e)$ for each $w \in W$. However, the problem is that $T(W) \in I^V$ is a *square-integrable* function from G to V, so it is not clear that it is meaningful to evaluate it at the identity in G. We thus need to take a digression to discuss continuity.

Suppose that $T \in \mathrm{Hom}_G(W, I^V)$. We now use the assumption that W is a finite-dimensional representation: since W is finite-dimensional, we see that the image $T(W)$ is a finite-dimensional subspace of I^V. Furthermore, $T(W)$ is an invariant subspace of I^V because T is an intertwining operator. Furthermore, I^V is a closed invariant subspace of the Hilbert space

$$L^2(G,V) = \left\{ f : G \to V \,\middle|\, ||f||^2 = \int_G ||f(g)||_V^2 < \infty \right\}.$$

with the usual action

$$L(g)f(x) = f(g^{-1}x) \in V$$

for $f \in L^2(G,V)$ and $x, g \in G$.

We then use the fact, which we prove in a lemma following this theorem, that finite-dimensional invariant subspaces of $L^2(G,V)$ consist of continuous functions from G to V. In particular, we see that $T(W)$ is a finite-dimensional subspace of I_c^V (that is, it consists of continuous functions). For each $w \in W$, we may thus set

$$S_T w = T(w)(e) \in V.$$

It is clear that $S_T : W \to V$ is linear. Because both V and W are finite-dimensional, continuity is immediate. Next we show that S_T is an H-intertwining operator:

$$\begin{aligned} S_T(\pi(h)w) &= T(\pi(h)w)(e) \\ &= [\mathrm{Ind}_H^G(h)(Tw)](e) \text{ (T intertwines π and Ind_H^G)} \\ &= (Tw)(h^{-1}) \\ &= \sigma(h)(Tw)(e) \text{ (equivariance property of I^V)} \\ &= \sigma(h)(S_T w). \end{aligned}$$

for all $w \in W$ and $h \in H$.

Finally, we must show that $SR = \mathrm{Id}$ and $RS = \mathrm{Id}$. We demonstrate that $SR = \mathrm{Id}$ and leave the other direction to the reader. If $T \in \mathrm{Hom}_H(V,W)$, then we recall that $(R_T w)(g) = T(\pi(g^{-1}w) \in V$ for all $w \in W$ and $g \in G$. Thus, $S_{R_T} w = (R_T w)(e) = Tw$ and we are done. □

Remark 3.1.2. We now consider another way of looking at the Frobenius Reciprocity theorem. Suppose that $(\rho, \mathcal{H})$ is a unitary representation (not necessarily irreducible) of G. Then one may decompose $(\rho, \mathcal{H})$ into a direct sum of irreducible representations:

$$\mathcal{H} \cong \bigoplus_{\pi \in \widehat{G}} m(\pi,\rho) V_\pi,$$

where each (π, V_π) is an irreducible representation of G, $m(\pi,\rho)$ is a nonnegative integer, and $m(\pi,\rho)V_\pi$ denotes a direct sum of $m(\pi,\rho)$ of V_π. We say that $m(\pi,\rho)$ is the **index of** π **in** ρ, and it measures the number of times that π appears as a subrepresentation of ρ.

We can use Schur's Lemma to provide a nice way to calculate this number. Suppose that (σ, V_σ) is an irreducible representation of G. Then we see that the dimension space of intertwining operators from V_σ to $\mathcal{H}$ is given by:

$$\begin{aligned}\dim \operatorname{Hom}_G(V_\sigma, \mathcal{H}) &= \sum_{\pi \in \widehat{G}} m(\pi, \rho) \dim \operatorname{Hom}_G(V_\sigma, V_\pi) \\ &= m(\sigma, \rho),\end{aligned}$$

where we used Schur's lemma to see that $\dim \operatorname{Hom}_G(V_\sigma, V_\pi) = 1$ if $\sigma = \pi$ and $\dim \operatorname{Hom}_G(V_\sigma, V_\pi) = o$ if $\pi \ncong \sigma$. We thus immediately arrive at the following corollary of Frobenius Reciprocity:

Corollary 3.1.3. *Let G be a compact group with a closed subgroup H, and suppose that (σ, V) and (π, W) are finite-dimensional unitary representations of H and G, respectively. Then*

$$m(\sigma, \operatorname{Res}_H^G \pi) = m(\pi, \operatorname{Ind}_H^G \sigma).$$

That is, the irreducible representation σ of H appears in $\operatorname{Res}_H^G \pi$ the number of times that the irreducible representation π of G appears in $\operatorname{Ind}_H^G \sigma$.

Example 3.1.4. As an example of the power of Frobenius Reciprocity, we show how one may use it to provide an alternate proof of the Peter-Weyl theorem on the decomposition of $L^2(G)$ for a compact group G.

We saw earlier that the regular representation $(L, L^2(G))$ is equivalent to the representation $\operatorname{Ind}_{\{e\}}^G 1_{\{e\}}$ induced from the trivial representation of the trivial subgroup. Now we consider an irreducible representation (π, V) of G. Notice that $\operatorname{Res}_{\{e\}}^G \pi$ is the trivial action on V, so that every subspace is $\{e\}$-invariant. That is, (π, V) decomposes into a direct sum of $\dim V$ copies of the trivial representation of the trivial subgroup $\{e\}$. Thus, by Frobenius Reciprocity, we have $\dim V = m(1_{\{e\}}, \operatorname{Res}_{\{e\}}^G \pi) = m(\pi, \operatorname{Ind}_{\{e\}}^G 1_{\{e\}})$. That is, exactly $\dim V$ copies of (π, V) of G appear in the decomposition of $L^2(G)$:

$$L^2(G) \cong_G \bigoplus_{\pi \in \widehat{G}} (\dim V_\pi) \cdot V_\pi$$

The exact statement of Peter-Weyl may be furthermore recovered by similar arguments.

Example 3.1.5. As a slightly more interesting example, we recall that if H is a closed subgroup of a compact group, then the quasi-regular representation $(L_{G/H}, L^2(G/H))$ is the same as the induced representation $\operatorname{Ind}_H^G 1_H$, where 1_H denotes the trivial representation of H. Thus we have by Frobenius Reciprocity that if (π, V) is an irreducible representation of G, then $m(\pi, L_{G/H}) = m(1_H, \pi)$. Now a subspace W of V corresponds to the trivial representation of H if it is one-dimensional and the action of H on W does nothing. In other words, we see that that $m(1_H, \pi) = \dim V^H$, where

$$V^H = \{v \in V | \pi(h)v = v \text{ for all } h \in H\}$$

is the space of H-fixed vectors in V. Hence we have the decomposition

$$L^2(G) \cong_G \bigoplus_{\pi \in \widehat{G}} (\dim(V_\pi)^H) \cdot V_\pi.$$

This reduces the analytic problem of understanding L^2-functions on G/H into the algebraic problem of finding H-invariant vectors in the (finite-dimensional) irreducible representations of G.

We end the section by sketching the proof of the lemma that we used in the proof of Frobenius Reciprocity.

Lemma 3.1.6. *If G is a compact group and V is a finite-dimensional complex vector space, then all finite-dimensional invariant subspaces of the regular representation $(L, L^2(G, V))$ is contained in the space of continuous functions from G to V.*

Proof sketch. It is enough to prove that finite-dimensional invariant subspaces of $L^2(G)$ consist of continuous functions: in fact, V is finite-dimensional and $L^2(G, V)$ decomposes into a direct sum of $\dim V$ copies of the regular representation $L^2(G)$.

Suppose now that W is a finite-dimensional invariant subspace of $L^2(G)$. Because G is compact, we see that $C(G) \subset L^2(G)$ (where $C(G)$ consists of the continuous functions). For any $f \in W$ and every continuous function $g \in C(G)$, we may thus consider the convolution $f * g$. One shows using Lebesgue dominated convergence that $f * g \in C(G) \subset L^2(G)$.

Furthermore, we see that

$$g * f(x) = \int_G g(y) f(y^{-1}x) dy = \int_G g(y)(L(y)f)(x) dy.$$

But we know that $g(y)L(y)f \in W$ for all $y \in G$ since W is an invariant subspace. One then shows by a limiting argument that $g * f = \int_G g(y)(L(y)f) dy \in W$ for all $g \in C(G)$ and $f \in W$.

Now if we could find a continuous function $h \in C(G)$ such that $h * f = f$ for each $f \in W$, then we would automatically have that $W \subseteq C(G)$. Unfortunately, it is widely known that the convolution product does not have an identity if G is not a discrete group. However, there do exist *approximate identities*. That is, it is possible to find a sequence $\{g_n\}_{n \in \mathbb{N}}$ of continuous functions on G such that $g_n * f \to f$ as $n \to \infty$ for all $f \in L^1(G)$.

In particular, if $f \in W$, then $g_n * f$ is in both W and $C(G)$ for all $n \in N$ and yet $g_n * f \to f$ in $L^1(G)$. Because W is a finite-dimensional space, there is only one Hausdorff topology, and thus in particular we see that $C(G) \cap W$ is a dense subspace of W and must therefore be equal to W. □

Remark 3.1.7. The precise construction of this approximate identity is not difficult but beyond the scope of these notes (see [6] for more on this topic). Essentially one chooses a sequence of continuous functions g_n which satisfy $\int_G g_n(x) dx = 1$ and become closer and closer to the "delta function": that is, they converge pointwise to zero everywhere except at the identity, where they approach ∞.

Remark 3.1.8. if G is a Lie group, then one may also construct approximate identities of smooth functions, so that in fact every finite-dimensional subspace of $L^2(G)$ consists of smooth functions.

3.2 Direct Sums and Induction in Stages

For compact groups, most of the important properties of induced representations follow quickly from Frobenius Reciprocity. However, that theorem is blatantly false for noncompact groups: take, for instance, the regular representation of $\mathbb{R}$ on $L^2(\mathbb{R})$. If Frobenius Reciprocity held for $\mathbb{R}$, then the regular representation would be equal to a direct sum of all the irreducible representations of $\mathbb{R}$. However, as we have already seen, $\mathbb{R}$ does not have any irreducible subrepresentations!

It is thus necessary to prove the most important properties of induced representations one by one. In this section we include two of the most important properties of induction: induction in stages and commutativity with direct sums.

Theorem 3.2.1. *If H is a closed subgroup of a locally compact group G and (π, V) and (σ, W) are unitary representations of H, then*

$$\operatorname{Ind}_H^G(\pi \oplus \sigma) \cong \operatorname{Ind}_H^G \pi \oplus \operatorname{Ind}_H^G \sigma.$$

In fact, one can prove a stronger result: if we have a measure space (μ, X) and a unitary representation $(\pi_x, \mathcal{H}_x)$ for each $x \in X$, then one can show that:

$$\operatorname{Ind}_H^G \left(\int_X^{\oplus} \pi_x d\mu(x) \right) \cong \int_X^{\oplus} \left(\operatorname{Ind}_H^G \pi_x \right) d\mu(x).$$

Proof. □

Theorem 3.2.2 (Induction in Stages)**.** *Suppose that K, H, G are locally compact groups with $K \leq H \leq G$, where each group is a closed subgroup of the next. If (σ, V) is a unitary representation of G, then $\operatorname{Ind}_H^G(\operatorname{Ind}_K^H \sigma) \cong \operatorname{Ind}_K^G \sigma$.*

Proof. □

3.3 Three Roads to Induction

3.4 Intertwining Operators Between Induced Representations

3.5 Imprimitivity

Bibliography

[1] D. Arnal, *Symmetric nonself-adjoint operators in an enveloping algebra.* J. Functional Analysis **21**, 432–447 (1975).

[2] T. Bröcker and T. tom Dieck, *Representations of Compact Lie Groups*, Springer 1985.

[3] G. van Dijk, *Introduction to Harmonic Analysis and Generalized Gelfand Pairs.* Walter de Gruyter, 2009.

[4] J. Dixmier, *C^*-Algebras.* North-Holland, 1977.

[5] J. Dixmier, P. Malliavin. *Factorisations de fonctions et de vecteurs indéfiniment différentiables.* Bull. Sci. Math. (2) **102** (1978), No. 4, 307–330.

[6] R. Fabec, G. Olafsson, *Non-Commutative Harmonic Analysis*, Drexville Publishing 2014.

[7] J. Faraut, *Analysis on Lie Groups: an Introduction*, Cambridge 2008.

[8] G. Folland, *A Course in Abstract Harmonic Analysis.* CRC, 1995.

[9] H. Glöckner, *Direct limit Lie groups and manifolds.* J. math. Kyoto Univ. **3** (2003), 1–26.

[10] R. Goodman and N. Wallach, *Symmetry, Representations, and Invariants.* Springer 2009.

[11] F. Greenleaf and M. Moskowitz, *Cyclic vectors for representations of locally compact groups*, Mathematische Annalen **190** (1971), 265–268.

[12] Harish-Chandra, *Representations of a semisimple Lie group on a Banach space. I*, Trans. Amer. Math. Soc. **75** (1953), 185–243.

[13] S. Helgason, *A duality for symmetric spaces with applications to group representations.* Adv. Math. **5** (1970), 1–154.

[14] —, *Differential Geometry, Lie Groups, and Symmetric Spaces.* Academic Press, 1978.

[15] —, *Groups and Geometric Analysis.* Academic Press, 1984.

[16] —, *Geometric Analysis on Symmetric Spaces*, Second Edition. AMS 2008.

[17] A.W. Knapp, *Lie Groups: Beyond an Introduction,* Second Edition. Birkhäuser 2002.

[18] R. Lipsman, *The Plancherel formula for the horocycle space and generalizations,* J. Australian Math. Soc. Series A **61** (1996), No. 1, 42–56.

[19] G. Mackey, *Harmonic analysis as the exploitation of symmetry,* Bull. Amer. Math. Soc. (N.S.) **3** (1980), No. 1, 543–698.

[20] —, *The Theory of Unitary Group Representations*, U. Chicago 1976.

[21] F. I. Mautner, *Geodesic flows on Riemannian symmetric spaces*, Ann. Math. **65** (1957), No. 3, 416–431.

[22] W. Rudin, *Functional Analysis*, Second Edition. McGraw-Hill 1991.

[23] Y. S. Samoilenko, *Spectral Theory of Families of Self-Adjoint Operators*. Springer 1991.