• **Interlude:** Calculus on spheres: invariant integrals, invariant $\Delta = \Delta^S$, integration-by-parts, etc.

Decomposition of $L^2(S^{n-1})$ into Δ^S-eigenfunctions.

Representation theory of orthogonal groups $O(n, \mathbb{R})$ or $SO(n, \mathbb{R})$.

... combine to prove

Hecke's identity: For a homogeneous, degree d *harmonic* polynomial P on $\mathbb{R}^n$, $P(x)\, e^{-\pi|x|^2}$ is a Fourier transform eigenfunction with eigenvalue i^{-d}:

$$\Big(P(x)\, e^{-\pi|x|^2}\Big)\widehat{\ }(\xi) \;=\; i^{-d} \cdot P(\xi)\, e^{-\pi|\xi|^2}$$

We need the *representation theory* of $SO(n, \mathbb{R})$

A little representation theory: Given a (topological) group G, a group homomorphism $T : G \to \mathrm{Aut}^o_{\mathbb{C}}(V)$ to the group of continuous $\mathbb{C}$ linear automorphisms of a complex vector space V is a **representation** of G (on V). The representation is **finite-dimensional** when V is. When the group homomorphism $G \to \mathrm{Aut}^o_{\mathbb{C}}(V)$ is understood, the standard abuse is to say that V *itself* is the representation, not naming the homomorphism. In that case, rather than writing $(Tg)(v)$ for $g \in G$ and $v \in V$, write simply gv for the action of g on v.

There is a *continuity* requirement, that

$$G \times V \longrightarrow V \quad \text{is } \textit{continuous}$$

For finite-dimensional V, there is a (provably) unique (topological vector space) topology, so need not be specified explicitly. Topologies on infinite-dimensional V must be specified. The topology on G should be clear from context.

A G-**subrepresentation** of a representation V of G is a complex vector subspace W of V stable under the action of G... and when V is infinite-dimensional W must be topologically *closed*.

A representation V of G is **irreducible** if there is no *proper* G-subrepresentation, that is, if there is no G-subrepresentation of V other than $\{0\}$ and V itself.

A G-homomorphism from one G-representation V to another G-representation W is a complex-linear map $\varphi : V \to W$ commuting with the action of G in the sense that

$$\varphi(gv) \;=\; g\,\varphi(v) \qquad\qquad (\text{for } g \in G,\ v \in V)$$

Such G-homomorphisms are also called G-morphisms, G-maps, or also G-**intertwinings**, or G-**intertwining operators**. The collection of all G-intertwinings of V to W is $\mathrm{Hom}_G(V, W)$.

Theorem: *(instance of Schur's Lemma)* For a finite-dimensional irreducible representation V of a group G, any G-intertwining $\varphi : V \to V$ of V to itself is *scalar*.

Proof: First, claim that the collection $\mathrm{Hom}_G(V, V)$ of all G-intertwinings of finite-dimensional V to itself is a *division ring*. Indeed, given $\varphi \neq 0$ in the ring $\mathrm{Hom}_G(V, V)$, the image $\varphi(V)$ is readily seen to be a G-subrepresentation of V. For φ not the zero map, since V is irreducible, the image is either $\{0\}$ or V, so must be V since $\varphi \neq 0$, and φ is *surjective*. Similarly, the kernel of φ is a G-subrepresentation, so is either the whole V, impossible since φ is not the zero map, or is $\{0\}$. Thus, φ is injective. Thus, φ is a bijection, and therefore has an inverse (easily seen to be a G-map). Thus, non-zero elements of the ring $\mathrm{Hom}_G(V, V)$ have multiplicative inverses, and $\mathrm{Hom}_G(V, V)$ is a division ring.

For V finite-dimensional the whole collection of complex-linear endomorphisms of V is finite-dimensional. Certainly $\mathbb{C}$ naturally lies inside the *center* of $\mathrm{Hom}_G(V, V)$. For $\varphi \in \mathrm{Hom}_G(V, V)$, the collection of rational expressions $\mathbb{C}(\varphi)$ is a field, and is finite-dimensional as a vector space over the copy of $\mathbb{C}$ lying in the center of $\mathrm{Hom}_G(V, V)$, so is *algebraic*. But $\mathbb{C}$ is algebraically closed (by Liouville's theorem), so $\varphi \in \mathbb{C}$. ///

Theorem: The space H_d of harmonic homogeneous total-degree d polynomials is an *irreducible* $SO(n, \mathbb{R})$-representation.

Proof: Suppose, to the contrary, that X is a proper $SO(n, \mathbb{R})$-subspace of H_d. Then the orthogonal complement Y of X inside H_d with respect to the $SO(n, \mathbb{R})$-invariant inner product is also $SO(n, \mathbb{R})$-stable.

The subspace X consists of continuous functions, so, for $x \in S^{n-1}$, the functional $f \to f(x)$ is a linear functional on X. By an especially easy case of Riesz-Fischer, there is $F_x \in X$ such that $f(x) = \langle f, F_x \rangle$ for all $f \in X$. Since X is rotation-invariant and not $\{0\}$, the functional $f \to f(x)$ cannot be 0 on all of X, so $F_x \neq 0 \in X$.

Similarly, there is $0 \neq \Phi_x \in Y$ such that $f(x) = \langle f, \Phi_x \rangle$ for all $f \in Y$.

By rotating, without loss of generality $x = e_n$. Then the following funny lemma proves that F_x and Φ_x must be scalar multiples of each other, contradiction. ///

Let $SO(n-1)$ be the smaller orthogonal group which is the *isotropy group* of the point $e_n = (0, \ldots, 0, 1)$.

Lemma: On $\mathbb{R}^n$, for each fixed d, in H_d there is a unique (up to constant multiples) $SO(n-1, \mathbb{R})$-fixed vector f, that is, so that $h \cdot f = f$ for every $h \in SO(n-1, \mathbb{R})$.

Proof: A function invariant under $SO(n-1)$ must be of the form

$$f_o(x) \;=\; F(\rho^2, t) \qquad \text{(where } \rho^2 = x_1^2 + \ldots + x_{n-1}^2 \text{ and } t = x_n\text{)}$$

Computing,

$$\begin{aligned}\Delta f_0 \;&=\; \textstyle\sum_i \dfrac{\partial}{\partial x_i}(2x_i F_1) + F_{22} \;=\; \textstyle\sum_i (2F_1 + 4x_i^2 F_{11}) + F_{22} \\ &=\; 2(n-1)F_1 + 4\rho^2 F_{11} + F_{22}\end{aligned}$$

where subscripts denote derivatives.

Write the function as a polynomial in t with coefficients functions of ρ. These coefficients are necessarily *homogeneous* functions of ρ, so are *powers* of ρ. Write these as powers of ρ^2.

$$f_o(x) = F(\rho^2, t)$$

$$= c_d t^d + c_{d-1} t^{d-1} (\rho^2)^{1/2} + c_{d-2} t^{d-2} (\rho^2)^1 + \ldots + c_o (\rho^2)^{d/2}$$

The harmonic-ness $0 = 2(n-1)F_1 + 4\rho^2 F_{11} + F_{22}$ rewritten in powers of t gives a recurrence for the coefficients c_i. Explicitly, looking at the $(i-2)^{th}$ power of t in the harmonic-ness condition,

...

$$\begin{aligned} 0 &= 2(n-1)\Big(\frac{\partial}{\partial(\rho^2)}\Big)\Big(c_{i-2}(\rho^2)^{(d-i+2)/2}\Big) \\ &\quad +4\rho^2\Big(\frac{\partial}{\partial(\rho^2)}\Big)^2\Big(c_{i-2}(\rho^2)^{(d-i+2)/2}\Big)i(i-1)c_i(\rho^2)^{(d-i)/2} \\ &= \Big[2(n-1)\frac{1}{2}(d-i+2)+4\cdot\frac{1}{2}(d-i+2)\frac{1}{2}(d-i)\Big](\rho^2)^{(d-i)/2}c_{i-2} \\ &\qquad +i(i-1)(\rho^2)^{(d-i)/2}c_i \end{aligned}$$

Thus,

$$i(i-1)c_i \;=\; -(d-i+2)(n-1+d-i)c_{i-2}$$

and the c_i are determined completely from c_1 and c_o.

On the other hand, looking at the t^{d-1} term in the harmonic-ness relation,

$$0 \;=\; [2(n-1)\frac{1}{2}(\rho^2)^{-1/2} + 4\rho^2\frac{1}{2}(-\frac{1}{2})(\rho^2)^{-3/2}]c_1 = [(n-1)-1]\rho$$

So unless $n = 2$ we have $c_1 = 0$, and all coefficients are determined by c_o, giving the desired uniqueness result.

The case $n \;=\; 2$ can be treated more directly, from the easily demonstrable fact that

$$H_d \;=\; \mathbb{C}\cdot(x+iy)^d \;\oplus\; \mathbb{C}\cdot(x-iy)^d \qquad\qquad (\text{on } \mathbb{R}^2)$$

This proves the uniqueness lemma, and the irreducibility of H_d as $SO(n,\mathbb{R})$ representation. ///

The irreducibility of H_d is the key point, but there are a few other small requirements before Hecke's identity is completed.

As expected, two G-representations V, W are *isomorphic* when there is a vector space isomorphism $V \to W$ that is a G-hom. When the vector spaces are infinite-dimensional, the map is required to be a *topological* vector space isomorphism, as expected.

Similar to Schur's Lemma:

Proposition: For *non-isomorphic* irreducible finite-dimensional G-representations V, W, the only G-hom $\varphi : V \to W$ is the zero map.

Proof: The kernel and image of φ are G-subrepresentations, so are either the whole space or $\{0\}$. The case that $\varphi : V \to W$ is a vector space isomorphism is excluded by the non-isomorphism assumption. ///

Proposition: *(Complete Reducibility)* A finite-dimensional G representation V with a G-invariant inner product $\langle , \rangle$ is a finite orthogonal direct sum of irreducible G-subrepresentations.

Remark: A representation space with a G-invariant inner product is said to be *unitary*.

Proof: Induction on dimension. For V irreducible, we are done. Otherwise, V has a proper subrepresentation W, which is a direct sum of irreducibles, by induction. The orthogonal complement $W^\perp$ of W is immediately G-stable, so is a G-subrepresentation, and *also* is a direct sum of irreducibles, by induction. ///

Remark: There is no straightforward infinite-dimensional analogue of the previous, even for Hilbert spaces. The orthogonality argument still succeeds, but induction fails.

The *idea* is that the $SO(n, \mathbb{R})$-map

$$\# : H_d \longrightarrow \mathbb{C}[x]^{\leq d} \approx H_d \oplus (\text{other irreducibles})$$

must map H_d to the other copy of H_d, and, by Schur's lemma, be a scalar.

This idea is correct, but has not quite been proven so far.

Specifically, although we easily showed that, for irreducibles σ, τ, the space of G-homs $\mathrm{Hom}_G(\sigma, \tau)$ is either 0 or $\mathbb{C}$ depending on whether $\sigma \not\approx \tau$ or $\sigma \approx \tau$, this does not instantly address the question of G-maps to *sums* of irreducibles.

It is best to examines some clarifying structure.

Isotypes and co-isotypes A direct sum of a number of copies of an irreducible π is denoted

$$m \cdot \pi \;=\; \underbrace{\pi \oplus \ldots \oplus \pi}_{m}$$

Given an irreducible π of G, we want to specify a G-sub V^π of a G-rep'n V such that any map $m \cdot \pi \to V$ *factors (uniquely) through* V^π, that is, $m \cdot \pi \to V^\pi \subset V$, the π-*isotype* of V.

Dually, the π-co-isotype V_π of V is the quotient of V such that any map $V \to m \cdot \pi$ factors through V_π: $V \to V_\pi \to m \cdot \pi$.

A priori, existence is unclear, but on categorical grounds they are unique up to unique isomorphism if they do exist at all.

For *unitary* representations, the kernel of the map to the co-isotype has an orthogonal complement, so the co-isotype is naturally isomorphic to a sub-object, ... but in general we should not expect this simplicity.

Happily, for finite-dimensional irreducibles π of *compact* G, there is a natural *projector* to the π-isotype.

It is not obvious, but the history of these issues does reasonably lead to the following. For a finite-dimensional irreducible π of compact G, the *character* χ_π of G is a function on G defined by

$$\chi_\pi(g) \;=\; \mathrm{trace}\,\pi(g)$$

Proposition: For any G-representation V, the map

$$v \;\longrightarrow\; \chi_\pi \cdot v \;=\; \int_G \chi_\pi(g)\; g\,v\; dg$$

is a G-hom *projecting* $V \;\to\; V^\pi$, where $\pi^\vee$ is the *contragredient (dual)* representation of π.