

Linear Transformations and Group Representations

These notes are intended to follow the “Groups, Fields, and Vector Spaces” notes from 2014-2015.

Eigenvectors and Eigenvalues

The terms “linear transformation of (or on) V ”, “linear operator on V ”, and “member of $\text{Hom}(V, V)$ ” will be used interchangeably.

Definitions

Above we defined the determinant of a linear transformation A on V , and (by doing this in a coordinate-free manner) showed that it is an intrinsic property of A , i.e., one that is independent of the choice of basis. Here we use the determinant to find some other intrinsic properties of A .

For a linear transformation A in a vector space V , an eigenvector is v is, by definition, a nonzero vector that satisfies $Av = \lambda v$ for some scalar (field element) λ . λ is called the eigenvalue for A associated with v . λ is allowed to be 0, but not v . Note that eigenvalues and eigenvectors are defined in a coordinate-free fashion, so they are intrinsic properties of A .

Typically, a linear transformation has a whole a *set* of eigenvalues λ_j and associated eigenvectors v_j , satisfying $Av_j = \lambda_j v_j$.

We will initially work in a finite-dimensional vector space. This allows us to see the algebraic structure clearly, but it also uses some tools that do not apply in the infinite-dimensional case. (And, without further assumptions, many of the results for finite dimensions do not hold in the infinite-dimensional case.) But later we will add one more piece of structure – the “inner product” – which (a) restricts the infinite-dimensional spaces we can consider, but (b) ensures that the key results for finite-dimensional spaces will apply. Essentially, this happens because the “inner product” yields a notion of distance, and the notion of distance allows us to restrict attention to vectors that can be well-approximated by vectors in a finite-dimensional space.

For a finite-dimensional vector space V , (say, of dimension n) we can find the eigenvalues of A by solving the “characteristic equation” of A , namely, $\det(zI - A) = 0$. (Here, I is the identity transformation on V). This works because if $\det(zI - A) = 0$ is solved by $z = \lambda$, then $\lambda I - A$ is an operator that transforms a basis set into a set whose span is of at most dimension $n - 1$. So some linear combination of the basis set (say, v) must be mapped to 0 by $\lambda I - A$. And if $(\lambda I - A)v = 0$, then $\lambda Iv = Av$, i.e., $\lambda v = Av$.

If the dimension of V is n , the characteristic equation is a polynomial of degree n , i.e.,

$$\det(zI - A) = z^n - a_1 z^{n-1} + a_{n-2} z^{n-2} - \dots + (-1)^n a_n.$$

The characteristic equation leads to intrinsic descriptions

Note that since we didn't use coordinates to define the determinant, the characteristic equation is an intrinsic property of A . This means that each of its roots – i.e., each of the eigenvalues – are intrinsic properties of A .

To think about: in the setup that V is a space of signals and $Hom(V, V)$ a space of linear transformations (e.g., from input to output), and we are studying a particular A in $Hom(V, V)$, wouldn't it be nice to know describe A in terms of these intrinsic properties?

Each of its coefficients of the characteristic equation are also intrinsic properties of A . The coefficient a_n is the determinant (set $z = 0$ in the above), but the other coefficients yield new things. If you work through the definition of the determinant and carry it out with coordinates, you will see that a_1 is the sum of the elements on the diagonal of A . This is known as the “trace” of A , $tr(A)$.

If the field k is algebraically closed, then solutions of the characteristic equation will exist, and the characteristic equation completely factors:

$\det(zI - A) = (z - \lambda_1)(z - \lambda_2) \cdots (z - \lambda_n)$. (This is why we choose $k = \mathbb{C}$, the field of complex numbers: \mathbb{C} is algebraically closed.)

Equating coefficients with the characteristic equation shows that $tr(A)$ is the sum of the eigenvalues, and $\det(A)$ is the product of the eigenvalues.

Because the trace is the sum of the eigenvalues, it has another important property that we will use below: $tr(AB) = tr(BA)$. To see this, write $C = BA$, so $AB = ACA^{-1}$.

$tr(AB) = tr(BA)$ is thus equivalent to $tr(ACA^{-1}) = tr(C)$. Three ways to see this: (i) recognize that ACA^{-1} is the same transformation as C , written in a different basis set. (ii) The trace is the highest-power term in the characteristic equation, and ACA^{-1} has the same characteristic equation as C : $\det(\lambda I - ACA^{-1}) = \det(A(\lambda I - C)A^{-1}) = \det(\lambda I - C)$. (iii) Any eigenvector/eigenvalue pair for C , say, (λ, v) is also an eigenvector/eigenvalue pair for ACA^{-1} , namely (λ, Av) (and vice-versa): $ACA^{-1}(Av) = ACv = A\lambda v = \lambda(Av)$. (These arguments only holds if A is invertible, but they are readily extended to the case when A is not.)

To emphasize: the eigenvalues and eigenvalues of A do not depend on the coordinates chosen for V – so they form a coordinate-independent description of A . (Of course to *communicate* the eigenvectors v_j , one typically does need to choose coordinates.)

Eigenvalues define subspaces

Eigenvectors corresponding to the same eigenvalue form a common subspace. To see this, suppose v and w are both eigenvectors of A with the same eigenvalue λ . Then any linear combination of v and w also is an eigenvector of A with the eigenvalue λ .

$$A(av + bw) = aAv + bAw = a\lambda v + b\lambda w = \lambda(av + bw).$$

So we can talk about the eigenspace associated with an eigenvalue λ , namely, the set of all eigenvectors. This forms a subspace of the original space V .

Conversely, eigenvectors corresponding to different eigenvalues lie in different subspaces. Suppose instead that v is an eigenvector of A with the eigenvalue λ , and that W is a subspace of V with a basis set of eigenvectors w_m all of whose eigenvalues λ_m are distinct from λ .

Then v cannot be in W . For if v were in W , then we could write $v = \sum a_m w_m$. On the one hand,

$$Av = \lambda v \text{ so } Av = \sum \lambda a_m w_m.$$

On the other hand, we could write $Av = A\left(\sum a_m w_m\right) = \sum a_m A(w_m) = \sum \lambda_m a_m w_m$. Since the w_m are a basis set, they are linearly independent, so the coefficients of the w_m must match in these two expansions of Av . That is, for each m , we would need to have $(\lambda - \lambda_m)a_m = 0$. Since we have assumed that for all m , $\lambda_m \neq \lambda$, it follows that all the a_m must be 0 – so v is not an eigenvector.

The above comment guarantees that eigenvectors corresponding to distinct eigenvalues are linearly independent.

While there is no guarantee, though, that the eigenvectors span V , there are many circumstances when this is the case. One case is that the characteristic equation has all distinct roots. Others are mentioned below.

When the eigenvectors form a basis

Say there is a special linear transformation T (specified by the problem at hand), with all of its eigenvalues λ_j distinct. Then its eigenvectors v_j form a basis that is singled out by T . It is also a basis in which the action of T on any vector $v \in V$ is simple to specify: since $v = \sum a_j v_j$ for some set of coefficients a_j , then $T(v) = T\left(\sum a_j v_j\right) = \sum T(a_j v_j) = \sum a_j T(v_j) = \sum a_j \lambda_j v_j$.

Another way of looking at this is that if you use the eigenvectors v_j as the basis set, then the

matrix representation of T is $T = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{pmatrix}.$

Note also that if the eigenvalues of T are λ_j , then the eigenvalues of aT are $a\lambda_j$, the eigenvalues of T^2 are λ_j^2 , etc., and, we can even interpret $f(T)$ as a transformation with eigenvalues $f(\lambda_j)$, for any function f .

Shared eigenvectors and commuting operators

Say A and B are linear transformations, and $AB = BA$. If v is an eigenvector of B with eigenvalue λ , then Av is also an eigenvector of B with eigenvalue λ . This is because

$$B(Av) = (BA)v = (AB)v = A(Bv) = A(\lambda v) = \lambda(Av).$$

Now further suppose that the eigenspace of B corresponding to eigenvalue λ has dimension 1 – for example, B is the T of the previous section, which has all of its eigenspaces of dimension 1. It follows that v is also an eigenvector of A . This is because (under the dimension-1 hypothesis) Av and v are both in the same one-dimensional eigenspace of B , so it must be that Av is a multiple of v , i.e., $Av = \mu v$, i.e., v is an eigenvector of A . Since each eigenvalue of B is also an eigenvalue of A , the basis set that diagonalizes B also diagonalizes A .

Thus, even if all of the eigenvalues of A are not distinct, the fact that it commutes with an operator T whose eigenvalues are distinct means that the eigenvectors for T form a natural basis for A .

How this applies to signals and systems

Before proceeding further with the abstract development, it is useful to see how what we already have is applied to signals and systems.

V is a vector space of functions of time. Linear transformations on V arise as filters, as input-output relations, as descriptors of spiking processes, etc. We want to find invariant descriptors for linear transformations on V , and, if possible, a preferred basis set.

Example: linear filters

The transformation $w = Lv$, with

$$w(t) = \int_0^{\infty} L(\tau)v(t-\tau)d\tau \tag{1}$$

is a linear transformation on V . View $v(t)$ as an input to a linear filter, $w(t)$ as an output. Here, $L(t)$, which describes L , is called the “impulse response”: $L(t)$ is the response $w = Lv$ when

$v(t) = \delta(t)$, the delta-function impulse, since $L\delta(t) = \int_0^{\infty} L(\tau)\delta(t-\tau)d\tau = L(t)$. (This is the basic property of the delta-function.)

Example: smoothing

Smoothing transformations are also linear transformations $w = Lv$, with

$$w(t) = \int_{-\infty}^{\infty} L(\tau)v(t-\tau)d\tau. \quad (2)$$

For example, take $L(t) = 1/(2h)$ if $|t| < h$, 0 otherwise -- “boxcar smoothing”. Or take $L(t)$ to be a Gaussian. L in this context is often called the “smoothing kernel.”

Below we show that time translation is an operator that commutes with the above L 's. We will then use this to determine a natural basis for V , in which it is simple to describe the action of the L 's, and to see how they combine.

Other examples that benefit from this setup will arise when we discuss point processes.

Time-translation invariance

In the above examples, the transformation L is “time-translation invariant” -- independent of absolute clock time. This crucial property can be formulated algebraically as a statement that certain operators commute.

To do this, we define the time-shift operator D_T on V as follows:

$$(D_T v)(t) = v(t+T). \quad (3)$$

That is, D_T advances time by T units. Note that this is a linear transformation.

This is equivalent to an expression of the form (2), if we allow L to be a generalized function, $L(\tau) = \delta(\tau+T)$.

($\delta(x)$ is a generalized function that satisfies $\int_{-\infty}^{\infty} \delta(x-a)g(x)dx = g(a)$. It can be thought of as

the limiting case of a “blip” of width Δ and height $1/\Delta$, as $\Delta \rightarrow 0$). The “limit” is not a limit in the usual sense, but integrals of the delta-function have a limit, which is all we need. And it is also in keeping with our policy that if something makes sense with arbitrarily fine discretizations of time, then there should be an extension of it that makes sense in the continuum limit.)

With $L(\tau) = \delta(\tau+T)$, eq. (2) becomes $\int_{-\infty}^{\infty} \delta(\tau+T)v(t-\tau)d\tau = v(t+T) = (D_T v)(t)$, since the only contribution to the integral is when the argument of the delta-function is zero, i.e., when $\tau+T=0$, i.e., $\tau=-T$.

Time-translation invariance of a linear operator A means that A has the same effect if the absolute clock time is unchanged. That is, $AD_T = D_T A$. The left-hand side means, first shift

absolute time and then apply A ; the right-hand side means, first apply A and then shift absolute time.

To show that operators defined by eq. (1) or (2) are time-translation invariant (it suffices to consider the case of eq. (2)), we need $LD_T v = D_T L v$:

$$(LD_T v)(t) = \int_{-\infty}^{\infty} L(\tau)(D_T v(t-\tau))d\tau = \int_{-\infty}^{\infty} L(\tau)(v(t-\tau+T))d\tau = \int_{-\infty}^{\infty} L(\tau'+T)v(t-\tau')d\tau'$$

where the last equality follows by substituting $\tau' = \tau - T$. Consequently,

$$\int_{-\infty}^{\infty} L(\tau'+T)v(t-\tau')d\tau' = \int_{-\infty}^{\infty} D_T L(\tau')v(t-\tau')d\tau' = (D_T L v)(t).$$

Since D_T itself (for any T) is of the form (2), this means that $D_{T'}D_T = D_T D_{T'}$, i.e., any two D_T 's commute. Thus, we should expect to find a set of vectors that are eigenvectors for all of the D_T 's. These will turn out to be all distinct, and to span V .

Then, these eigenvectors must also be eigenvectors for any time-translation invariant operator L (i.e., an L for which $LD_T = D_T L$). Expressed in this basis, L (including all transformations of the form (1) or (2)) are diagonal – and thus, easy to manipulate.

We will first find the eigenvectors and eigenvalues of D_T “by hand.” Then, we see how their properties arise because of the way that the time-translation group acts on the domain of the functions of V .

What are the eigenvectors and eigenvalues of D_T ?

Let's find the vectors v that are simultaneous eigenvectors of all the D_T 's.

First, observe that $D_S(D_T v)(t) = v(t+T+S) = D_{T+S}(v)$, so that $D_S D_T = D_{T+S}$. Intuitively, translating in time by T , and then by S , is the same as translating in time by $T+S$. Abstractly, the mapping $W : T \rightarrow D_T$ is a homomorphism of groups. It maps elements T of the group of the real numbers under addition (time translation) to some isomorphisms D_T of V .

Say $v(t)$ is an eigenvector for all of the D_T 's. We next see how the eigenvalue corresponding $v(t)$ depends on T . Say the eigenvalue associated with $v(t)$ for D_T is $\lambda(T)$. Since $D_S D_T = D_{T+S}$, $v(t)$ is an eigenvector of D_{T+S} , with eigenvalue $\lambda(T+S) = \lambda(T)\lambda(S)$. So the dependence of the eigenvalue on T must satisfy $\lambda(T+S) = \lambda(T)\lambda(S)$. Equivalently, $\log \lambda(T+S) = \log \lambda(T) + \log \lambda(S)$. That is, $\log \lambda(T)$ must be proportional to T . Choose a proportionality constant c . $\log \lambda(T) = cT$ implies that $\lambda(T) = e^{cT}$, for some constant c .

This determines $v(t)$: This is because $v(t+T) = (D_T v)(t) = \lambda(T)v(t) = e^{cT}v(t)$.

Choosing $t = 0$ now yields $v(T) = v(0)e^{cT}$, so these are the candidates for the simultaneous eigenvectors of all of the D_T 's.

If we choose a value of c that has a positive real part, then $v(T)$ gets infinitely large as $T \rightarrow \infty$. But if we choose a value of c that has negative real part, then $v(T)$ gets infinitely large as $T \rightarrow -\infty$. So the only way that we can keep $v(T)$ bounded for all T is to choose c to be pure imaginary. With $c = i\omega$, $v(T) = e^{i\omega T}$.

The above elementary calculation found all the eigenvalues and eigenvectors of the translation operator, but it did not guarantee that the eigenvectors span the space (i.e., form a basis for it). This is also true, and it follows from some very general results about how groups (in this case, the translation group) act on vector spaces (in this case, functions on the line). We'll get a look at this general result below.

Thus, the set of $v(T) = e^{i\omega T}$ (for all ω) not only form the complete set of eigenvectors of each of the D_T 's, but also form a basis for a vector space of complex-valued functions of time. They thus constitute natural coordinates for this vector space, in which time-translation-invariant linear operators are all diagonal. Fourier analysis is simply the re-expression of functions of time in these coordinates. This is also why Fourier analysis is useful. Because linear operators are diagonal when expressed in these new coordinates, the actions of filters can be carried out by coordinate-by-coordinate multiplication, rather than integrals (such as eq. (1)).

Hilbert spaces

To get started with this general result, we need to add one more piece of structure to vector spaces: the inner product. An inner product (or "dot-product"), essentially, adds the notion of distance. A vector space with an inner product, and in which all inner products are finite, is known as a Hilbert space. In a Hilbert space, it is possible to make general statements about what kinds of linear transformations have a set of eigenvectors that form a basis.

Some preliminary comments:

We can always make a finite-dimensional vector space into a Hilbert space, since we are guaranteed a basis (choose a set of coordinate axes), and we can choose the standard dot-product in that basis. This determines a notion of distance, and hence, which vectors are "unit vectors", i.e., what are the spheres. Had we chosen a different set of coordinate axes, say, ones that are oblique (in the basis of the first set), then we would have defined a different dot-product. But we could always find a linear transformation from the vector space to itself that transforms one dot-product into the other – this is the linear transformation that changes the first basis set into the second one. It would turn spheres into ellipsoids, and vice-versa. Thus, while adding Hilbert space structure to a finite-dimensional vector space does add a notion of "geometry", it doesn't allow us to prove things that we couldn't prove before – since Hilbert space structure was guaranteed, and all we are doing is choosing one example from an infinite set of possibilities.

The situation is very different for infinite-dimensional vector spaces, such as function spaces. Here, when we add a dot-product (and insist that it has a finite value), we actually need to exclude some functions from the space. As in the finite-dimensional case, adding the dot-product gives a notion of “geometry.” But it does something even more important: by excluding some functions from the vector space, it allows many of our intuitions from finite-dimensional vector spaces to generalize.

Definition of an inner product

An inner product (or “dot-product”) on a vector space V over the reals or complex numbers is a function from pairs of vectors to the base field, typically denoted $\langle v, w \rangle$ or $v \cdot w$. It must satisfy the following properties (where a is an element of the base field):

Symmetry: $\langle v, w \rangle = \langle w, v \rangle$ for $k = \mathbb{R}$, and $\langle v, w \rangle = \overline{\langle w, v \rangle}$ for $k = \mathbb{C}$. (Here and below, we denote the complex conjugate of a field element a by \bar{a}).

Linearity: $\langle av_1 + bv_2, w \rangle = a\langle v_1, w \rangle + b\langle v_2, w \rangle$, and $\langle v, aw_1 + bw_2 \rangle = \bar{a}\langle v, w_1 \rangle + \bar{b}\langle v, w_2 \rangle$. The second equality follows from the first one by applying symmetry.

Positive-definiteness: $\langle v, v \rangle \geq 0$ and $\langle v, v \rangle = 0$ only for $v = 0$.

The quantity $\langle v, v \rangle = \|v\|^2$ can be regarded as the square of the size of v , i.e., the square of its distance from the origin.

Implicit in the above definition is that $\langle v, w \rangle$ is finite. This does not mean that there is a universal upper limit to it that applies to all members of V , just that for any pair of vectors, $\langle v, w \rangle$ is a finite number.

Note that $\langle av, bw \rangle = a\bar{b}\langle v, w \rangle$. The necessity for the complex-conjugation is apparent if one considers $\langle iv, iv \rangle$. With complex-conjugation of the “ b ”, we find $\langle iv, iv \rangle = \bar{i}i\langle v, v \rangle = i(-i)\langle v, v \rangle = \langle v, v \rangle$, which is “good” – multiplication of v by a unit (i) does not change its length. But without complex conjugation, we’d find that $\langle iv, iv \rangle$ would equal $-\langle v, v \rangle$, i.e., positive-definiteness would be violated.

If $\langle v, w \rangle = 0$, v and w are said to be orthogonal.

The inner product, distances, triangle inequality, Cauchy-Schwartz, and angles

The quantity specified by $d(v, w) = \|v - w\| = \sqrt{\langle v - w, v - w \rangle}$ qualifies as a “metric” (i.e., a distance), because it is (a) symmetric, (b) non-negative, and (c) satisfies the triangle inequality $d(u, w) \leq d(u, v) + d(v, w)$. But demonstrating the triangle inequality is slightly harder than one might guess. To show that the triangle inequality follows from the positive-definiteness of the inner product: We want to show that $d(v, w) \leq d(v, x) + d(x, w)$, i.e.,

$\|v - w\| \leq \|v - x\| + \|x - w\|$; with $y = v - x$ and $z = x - w$ this is equivalent to

$$\|y + z\| \leq \|y\| + \|z\|, \text{ and to } \|y + z\|^2 \leq (\|y\| + \|z\|)^2 = \|y\|^2 + \|z\|^2 + 2\|y\|\|z\|.$$

Since $\|y\|^2 = \langle y, y \rangle$, $\|z\|^2 = \langle z, z \rangle$, and $\|y + z\|^2 = \langle y + z, y + z \rangle = \langle y, y \rangle + \langle z, z \rangle + \langle y, z \rangle + \langle z, y \rangle$, the latter is equivalent to $\langle y, z \rangle + \langle z, y \rangle \leq 2\|y\|\|z\|$, which is implied by

$|\langle y, z \rangle + \langle z, y \rangle|^2 \leq 4\|y\|^2\|z\|^2$. Since $\langle y, z \rangle + \langle z, y \rangle = \langle y, z \rangle + \overline{\langle y, z \rangle} = 2\operatorname{Re}\langle y, z \rangle \leq 2|\langle y, z \rangle|$, it suffices to show that $|\langle y, z \rangle| \leq \|y\|\|z\|$.

This final inequality is a form of the “Cauchy-Schwartz” inequality, which allows us to derive a notion of angles from the dot-product. The Cauchy-Schwartz inequality follows applying the positive-definiteness property to the difference between y and its projection on z (see section below on projections). We first find the projection of y on z : $p = z \frac{\langle y, z \rangle}{\langle z, z \rangle}$. The difference

between this and y is given by $q = y - z \frac{\langle y, z \rangle}{\langle z, z \rangle}$. So $y = p + q$, with p proportional to z and p

orthogonal to q (since $\langle q, p \rangle = \left\langle y - z \frac{\langle y, z \rangle}{\langle z, z \rangle}, z \frac{\langle y, z \rangle}{\langle z, z \rangle} \right\rangle = \langle y, z \rangle \frac{\langle y, z \rangle}{\langle z, z \rangle} - \langle z, z \rangle \left(\frac{\langle y, z \rangle}{\langle z, z \rangle} \right)^2 = 0$). That

is, y is the hypotenuse of a right triangle, p lying along z , and q being the perpendicular to the line of z . Now we can calculate:

$$\begin{aligned} \langle y, y \rangle &= \langle p + q, p + q \rangle = \langle p, p \rangle + \langle q, q \rangle = \left\langle z \frac{\langle y, z \rangle}{\langle z, z \rangle}, z \frac{\langle y, z \rangle}{\langle z, z \rangle} \right\rangle + \langle q, q \rangle \\ &= \langle z, z \rangle \left| \frac{\langle y, z \rangle}{\langle z, z \rangle} \right|^2 + \langle q, q \rangle = \frac{|\langle y, z \rangle|^2}{\langle z, z \rangle} + \langle q, q \rangle \end{aligned}$$

From this (and the non-negativity of $\langle q, q \rangle$), it follows that $\langle y, y \rangle \geq \frac{|\langle y, z \rangle|^2}{\langle z, z \rangle}$, which (since $\langle z, z \rangle$

is non-negative – if it were zero, the Cauchy-Schwartz inequality would have been trivial) implies the desired inequality $\langle y, y \rangle \langle z, z \rangle \geq |\langle y, z \rangle|^2$.

The Cauchy-Schwartz inequality enables us to interpret the quantity $\frac{|\langle y, z \rangle|}{\sqrt{\langle y, y \rangle \langle z, z \rangle}}$ as the cosine of the angle between the vectors y and z , as it is zero if the vectors are orthogonal, and has a maximal value of 1 if y and z are proportional to each other. The Cauchy-Schwartz inequality guarantees that it is always a real number in the range $[0,1]$.

Examples of the inner product

For a vector space of n -tuples of complex numbers, the standard inner product is

$$\langle u, v \rangle = \sum_{n=1}^N u_n \bar{v}_n. \quad (4)$$

Note that although we used coordinates to define these inner product, defining an inner product is not the same as specifying coordinates. As we will see below, we can choose alternate sets of coordinates that lead to exactly the same inner product. This is because the inner product only fixes a notion of distance, while the coordinates specify individual directions.

For functions of time, the standard inner product is

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(t) \bar{g}(t) dt. \quad (5)$$

We cannot consider all functions of time to form a Hilbert space with the inner product given by eq. (5), since this is not guaranteed to be finite. However, we can take V to be all functions of time for which the integral for $\langle f, f \rangle$ exists and is finite, i.e., that

$$\|f\|^2 = \langle f, f \rangle = \int_{-\infty}^{\infty} f(t) \bar{f}(t) dt = \int_{-\infty}^{\infty} |f(t)|^2 dt \text{ is finite. This guarantees (not obvious – Cauchy's}$$

inequality) that eq. (5) is finite as well, and makes this space (the “square-integrable functions of time”) to be a Hilbert space. It’s easy to find an example of a function that is perfectly well-defined, but for which the above integral does not exist – for example, a function that has any constant but nonzero value.

The inner product and the dual

An inner product is equivalent to specifying a correspondence between a vector space V and its dual V^* . (Remember, this was guaranteed in the finite-dimensional case, since the dimension of V and its dual are the same, but it is not guaranteed for the infinite-dimensional case.) That is, for each element v in V , the inner product provides a member φ_v of V^* , whose action is defined by $\varphi_v(u) = \langle u, v \rangle$. This correspondence is conjugate-linear (not linear), because $\varphi_{av} = \bar{a}\varphi_v$.

Some special kinds of linear operators

In a manner somewhat analogous to the above mapping between vectors and their duals, the inner product also specifies a mapping from an operator to its “adjoint”: the adjoint of an operator A is the operator A^* (sometimes written A^\dagger) for which $\langle Au, v \rangle = \langle u, A^*v \rangle$, for all u and v .

Two basic properties. First, the adjoint of the adjoint is the original operator.

$(A^*)^* = A$: Since, by definition, $(A^*)^*$ is defined as the operator for which $\langle A^*u, v \rangle = \langle u, (A^*)^*v \rangle$, we need to show $\langle A^*u, v \rangle = \langle u, Av \rangle$. This holds because $\langle A^*u, v \rangle = \overline{\langle v, A^*u \rangle} = \overline{\langle Av, u \rangle} = \langle u, Av \rangle$. (The middle equality is the definition of the adjoint, the first and third equalities are the conjugate-symmetry of the inner product.)

Second, the adjoint of a product is the product of the adjoints, in reverse order.

$B^*A^* = (AB)^*$, since $\langle u, B^*A^*v \rangle = \langle Bu, A^*v \rangle = \langle ABu, v \rangle$.

Third, the adjoint of the inverse is the inverse of the adjoint. That is, $(A^{-1})^* = (A^*)^{-1}$, since taking $B = A^{-1}$ in the above yields $(A^{-1})^*A^* = (AA^{-1})^* = I$, so $(A^{-1})^*$ is the inverse of A^* .

To see what the adjoint means in terms of coordinates: We write out $\langle Au, v \rangle = \langle u, A^*v \rangle$ and $\langle Au, v \rangle = \langle u, A^*v \rangle$, and force them to be equal, and look at the consequences for A and A^* .

Choose e_k (the vectors whose coordinates have a 1 in the unit position, and 0 elsewhere) as the basis elements and writes A , u and v in coordinates, i.e., $u = \sum_k u_k e_k$ and $v = \sum_k v_k e_k$, so that,

$$Au = \sum_j \left(\sum_k A_{jk} u_k \right) e_j, \text{ and } \langle Au, v \rangle = \sum_j \left(\sum_k A_{jk} u_k \right) \bar{v}_j = \sum_k \left(\sum_j A_{kj} u_j \right) \bar{v}_k.$$

Similarly, $A^*v = \sum_j \left(\sum_k A_{jk}^* v_k \right) e_j$ and $\langle u, A^*v \rangle = \sum_j \left(\sum_k \overline{A_{jk}^* v_k} \right) u_j$. Since this must be true for

all u_j and v_k , it follows that $A_{kj} \bar{v}_k = \overline{A_{jk}^* v_k}$, i.e., that $A_{kj} = \overline{A_{jk}^*}$. Thus, in coordinates, to find the adjoint, you (a) transpose the matrix (exchange rows with columns), and (b) take the complex-conjugate of its entries. And in the case of $k = \mathbb{R}$, the adjoint is the same as the transpose.

For the translation operator D_T acting on functions of the line, $(D_T v)(t) = v(t+T)$, the adjoint is

$(D_T)^* = D_{-T}$, i.e., $(D_{-T} v)(t) = v(t-T)$, since

$$\langle D_T u, v \rangle = \int u(t+T) \overline{v(t)} dt = \int u(t') \overline{v(t'-T)} dt = \langle u, D_{-T} v \rangle, \text{ where we've made the substitution } t' = t+T.$$

The adjoint allows us to define several special kinds of operators. These classes are intrinsic properties of $\text{Hom}(V, V)$ for a Hilbert space, i.e., they are defined in a coordinate-free manner but do require the specification of the inner product.

Self-adjoint operators

A “self-adjoint” operator A is an operator for which $A = A^*$. Self-adjoint operators have real eigenvalues, and, to some extent, can be thought of as analogous to real numbers. The fact that self-adjoint operators have real eigenvalues follows from noting that if $Av = \lambda v$, then

$$\lambda \langle v, v \rangle = \langle \lambda v, v \rangle = \langle Av, v \rangle = \langle v, A^* v \rangle = \langle v, Av \rangle = \langle v, \lambda v \rangle = \bar{\lambda} \langle v, v \rangle, \text{ so } \lambda = \bar{\lambda}.$$

For self-adjoint operators, eigenvectors with different eigenvalues are orthogonal.

Say $Av = \lambda v$ and $Aw = \mu w$, with $\lambda \neq \mu$. Then

$$\lambda \langle v, w \rangle = \langle \lambda v, w \rangle = \langle Av, w \rangle = \langle v, A^* w \rangle = \langle v, Aw \rangle = \langle v, \mu w \rangle = \bar{\mu} \langle v, w \rangle, \text{ so } \lambda = \bar{\mu} \text{ or } \langle v, w \rangle = 0. \text{ Since both } \lambda \text{ and } \mu \text{ are real, and they are assumed to be unequal, it follows that } \langle v, w \rangle = 0.$$

Unitary operators

A “unitary” operator A is an operator for which $AA^* = A^*A = I$, i.e., their adjoint is equal to their inverse. Unitary operators have eigenvalues whose magnitude is 1, and, to some extent, can be thought of as analogous to rotations, or to complex numbers of magnitude 1. The fact that unitary operators have eigenvalues of magnitude 1 follows from noting that if $Av = \lambda v$, then

$$|\lambda|^2 \langle v, v \rangle = \lambda \bar{\lambda} \langle v, v \rangle = \langle \lambda v, \lambda v \rangle = \langle Av, Av \rangle = \langle v, A^* Av \rangle = \langle v, v \rangle, \text{ so } |\lambda|^2 = 1.$$

If the base field is \mathbb{R} , then a unitary operator is also called an orthogonal operator.

For unitary operators, eigenvectors with different eigenvalues are orthogonal.

Say $Av = \lambda v$ and $Aw = \mu w$, with $\lambda \neq \mu$. Then

$$\langle v, w \rangle = \langle Av, Aw \rangle = \langle \lambda v, \mu w \rangle = \lambda \bar{\mu} \langle v, w \rangle = \frac{\lambda}{\mu} \langle v, w \rangle \text{ (with the last equality because } \mu \bar{\mu} = |\mu|^2 = 1). \text{ So if } \lambda \neq \mu, \text{ then } \langle v, w \rangle = 0. \text{ Conversely, a self-adjoint operator that has an inverse, and for which all eigenvalues have magnitude unity, is necessarily unitary.}$$

Note that the time-translation operator D_T is unitary, since its adjoint is D_{-T} , which is also its inverse.

Note also that the unitary operators in $\text{Hom}(V, V)$ form a group. It is closed under multiplication since $((AB)^*)^{-1} = (B^*A^*)^{-1} = (A^*)^{-1}(B^*)^{-1} = AB$ (if A and B have the property that their adjoint

is their inverse, then so does AB). Inverses are present because $A^{**} = A$, so $(A^*)^* A^* = AA^* = I$.

Projection operators

A “projection” operator is a self-adjoint operator P for which $P^2 = P$. One can think of P as a (geometric) projection onto a subspace – the subspace that is the range of P . It is also natural to consider the complementary projection, $Q = I - P$, as the projection onto the perpendicular (orthogonal) subspace. To see that Q is a projection, note

$$Q^2 = (I - P)^2 = (I - P)(I - P) = I - IP - PI + P^2 = I - P - P + P = I - P = Q. \text{ Also } PQ = P(I - P) = P - P^2 = 0. \text{ Also, the eigenvalues of a projection operator must be 0 or 1.}$$

This is because if $Pv = \lambda v$, then $Pv = P^2v = P(Pv) = P(\lambda v) = \lambda^2 v$ also, so $\lambda^2 = \lambda$, which solves only for 0 or 1.

A vector can be decomposed into a component that is in the range of P , and a component that is in the range of Q , and these components are orthogonal.

$v = Iv = (P + Q)v = Pv + Qv$, and $\langle Pv, Qv \rangle = \langle v, PQv \rangle = 0$ -- justifying the interpretation of P and Q as projections onto orthogonal subspaces.

Projections onto one-dimensional subspaces are easy to write. The projection onto the subspace determined by a vector u is the operator $P_u(v) = u \frac{\langle v, u \rangle}{\langle u, u \rangle}$.

To see that P_u is self-adjoint, note that $\langle P_u(v), w \rangle = \langle u, w \rangle \frac{\langle v, u \rangle}{\langle u, u \rangle} = \frac{\langle v, u \rangle \overline{\langle w, u \rangle}}{\langle u, u \rangle}$ but also

$\langle v, P_u(w) \rangle = \left\langle v, u \frac{\langle w, u \rangle}{\langle u, u \rangle} \right\rangle = \frac{\langle v, u \rangle \overline{\langle w, u \rangle}}{\langle u, u \rangle}$, where the last equality follows because the denominator must be real.

To see that $P_u^2 = P_u$, calculate

$$P_u^2(v) = u \frac{\langle P_u(v), u \rangle}{\langle u, u \rangle} = u \frac{\left\langle u \frac{\langle v, u \rangle}{\langle u, u \rangle}, u \right\rangle}{\langle u, u \rangle} = u \frac{\frac{\langle v, u \rangle}{\langle u, u \rangle} \langle u, u \rangle}{\langle u, u \rangle} = u \frac{\langle v, u \rangle}{\langle u, u \rangle}.$$

This construction can be extended to find projections onto multidimensional subspaces, specified by the range of an operator B . This is the heart of linear regression, and it will be useful for principal components analysis. Assuming that B^*B has an inverse, the projection can be written:

$P_B = B(B^*B)^{-1}B^*$. There’s an important piece of fine print here, in that the inverse of B^*B is only computed within the range of B .

A comment on how this definition of projection corresponds to the “usual” notion of projection as it applies to images. In the imaging context, one might consider a projection of an image $I(x, y, z)$ onto, say, the (X, Z) -plane by taking an average over all y . In our terms, this is a mapping from a function on a 3-d array of pixels (x, y, z) , to a function on a 2-d array of pixels (x, z) , i.e., a mapping between two vector spaces – and hence, might seem not to be a projection. But it is, in fact, a projection in our sense too. The functions on the 2-d array of pixels can also be regarded as functions on a 3-d array, but with exactly the same image “slice” for each value of y . The “projection” in our terms is to map $I = I(x, y, z)$ to PI , where the projection is defined

$$\text{by } (PI)(x, y, z) = \frac{1}{N_y} \sum_{y'=1}^{N_y} I(x, y', z).$$

Normal operators

A “normal” operator is an operator that commutes with its adjoint. Self-adjoint and unitary operators are normal. The only normal operators we will deal with here are either self-adjoint or unitary.

Idempotent operators

An “idempotent” operator is one whose square is itself, i.e., $A^2 = A$. It follows that all eigenvalues of an idempotent operator are 0 or 1, just like for a projection – but operators that are idempotent need not be projections.

Spectral theorem

Statement of theorem: in a Hilbert space, the eigenvectors of a normal operator form a basis. More specifically, the operator A can be written as

$$A = \sum_{\lambda} \lambda P_{\lambda} \tag{6}$$

where P_{λ} is the projection onto the subspace spanned by the eigenvectors of A with eigenvalue λ .

So this guarantees that the eigenvectors $v(t) = e^{i\omega t}$ of D_T form a basis, since D_T is unitary (and therefore, normal). It also tells us why we shouldn’t consider (possible) eigenvectors like $v(t) = e^{ct}$ for real c , since they are not in the Hilbert space. It also tells us that we can decompose vectors by their projections onto $v(t) = e^{i\omega t}$ (since they form a basis), and why representing operators in this basis (eq. (6)) results in a simple description of their actions.

But was it “luck” that D_T turned out to be unitary? Was it “luck” that, when the full set of operators was considered together, they had a common set of eigenvectors $v(t) = e^{i\omega t}$, and that there was one for each eigenvalue? Short answer: no, this is because the operators D_T expressed a symmetry of the problem.

The spectral theorem will also help us in another context, matters related to principal components analysis, which hinges on self-adjoint operators. In contrast to time series analysis (and its generalizations) in which unitary operators arise from *a priori* symmetry considerations, in principal components analysis, self-adjoint operators arise from the data itself.

Group representations

To understand why operators that express symmetries are unitary, and why they have common eigenvectors, and why they (often) have eigenspaces of dimension 1, we need to take a look at “group representation theory.” The basic setup is that vector spaces are often functions on a set of points, and if a group acts on a set of points, this induces transformations of the vector space. The transformations in the vector space “represent” the group. It turns out to be not too hard to find all possible representations of a group, and to write them in terms of “prime” (irreducible) representations. An (almost) elementary argument will show that because these representations are “prime”, they lead to a way to divide up the vector space, so that in each piece, the group acts in a simple way.

Unitary representations: definition and simple example

A unitary representation U of a group G is a structure-preserving mapping from the group to $Hom(V, V)$. More precisely, it is a group isomorphism from elements g of G into unitary operators U_g in $Hom(V, V)$.

Note that since U is a group isomorphism, $U_{gh} = U_g U_h$, where gh on the left is interpreted as multiplication in G , and $U_g U_h$ on the right is interpreted as composition in $Hom(V, V)$.

It’s worth looking at examples of group representations, since it makes it more impressive to find out that we can write out *all* the representations of a group. (The examples below don’t show this; they just show examples of the variety of representations that are possible.)

Example: cyclic groups

Consider the group \mathbb{Z}_n of addition (mod n), and let $V = \mathbb{C}$, i.e., V is the one-dimensional vector space of the complex numbers over itself. Then $U_p = e^{\frac{2\pi i}{n} p}$ is a representation of \mathbb{Z}_n . To check that it is an isomorphism, note that $U_p U_q = e^{\frac{2\pi i}{n} p} e^{\frac{2\pi i}{n} q} = e^{\frac{2\pi i}{n} (p+q)} = U_{p+q}$.

Example: the translation group on the line

Consider (again) time-shifts D_T acting on functions on the line by $(D_T v)(t) = v(t + T)$. This is a unitary representation, of the group of shifts on a line, in the vector space V of functions on the line.

Example: the dihedral group

Recall that the dihedral group D_n is the set of rotations and reflections of a regular n -gon. We can write out each of these rotations as a 2-d matrix, and obtain a 2-dimensional representation of the group.

Example: permutation groups

If a group that is presented as a permutation group, we can form unitary representations another way. We can write these permutations as permutation matrices, i.e., matrices that are mostly 0's, with a 1 in position (j,k) if element in position j is moved to position k . This is a representation too, as composing the permutations is equivalent to composing the matrices.

For the above cases (dihedral groups and permutation groups), we not have to check that the representations were in terms of unitary operators. This is because we were dealing with a finite group. In a finite group, every element g has an "order" m , i.e., a least positive integer for which $g^m = e$. If some operator L_g represents g , then (since the group representation preserves structure) $(L_g)^m = I$. Immediately, this means that any eigenvalue of L_g must satisfy $\lambda^m = 1$, which is a necessary condition for the operator to be unitary. But it is not sufficient, since (as homeworks showed), linear operators need not have a full set of eigenvectors. But from $(L_g)^m = I$, this possibility can be excluded (via consideration of the "Jordan Normal Form"), so that L_g is in fact guaranteed to be unitary. This approach fails for infinite groups, since we have no guarantee that each element has a finite order. But for infinite groups, the Hilbert space structure allows us to focus on unitary operators.

Example: the trivial representation

Finally, there is always the "trivial" representation, that takes every group element to the identity map on $\text{Hom}(V,V)$.

The character

The character χ_L of a representation L is a function from the group to the field. It is defined in terms of the trace: $\chi_L(g) = \text{tr}(L_g)$.

As noted above, the trace is the sum of the eigenvalues, and, $\text{tr}(AB) = \text{tr}(BA)$. In coordinates, this means that the trace is the sum of the diagonal elements. A simple consequence of this is that the character is invariant with respect to inner automorphisms: $\chi_L(hgh^{-1}) = \chi_L(g)$. To see this, $\chi_L(hgh^{-1}) = \text{tr}(L_{hgh^{-1}}) = \text{tr}(L_h L_g L_h^{-1}) = \text{tr}(L_h L_g (L_h)^{-1}) = \text{tr}(L_g) = \chi_L(g)$, where we have used the fact that L preserves structure, and that $\text{tr}(AB) = \text{tr}(BA)$.

Another way of putting this is that the character is constant on every “conjugate class” – the “conjugate class” of g is, by definition, the group elements of the form hgh^{-1} .

A couple of easy facts about characters:

The character of the trivial representation on a vector space V is equal to the dimension of V , since every group element is represented by the identity in V .

The character of any representation at the identity element is the dimension of the representation, since the representation at the identity element is the identity matrix.

For the translation group on the line, the character of the nontrivial *irreducible* representations will turn out to be the Fourier coefficients.

Combining representations

Two representations of the same group, U_1 in V_1 and U_2 in V_2 , can be combined to make a composite representation in $V_1 \oplus V_2$. A group element g maps to $U_{1,g} \oplus U_{2,g}$, where $U_{1,g} \oplus U_{2,g}$ acts on $v_1 \oplus v_2$ in the obvious way: $(U_{1,g} \oplus U_{2,g})(v_1 \oplus v_2) = U_{1,g}(v_1) \oplus U_{2,g}(v_2)$.

We can also define a group representation on $V_1 \otimes V_2$ in the same way:

$$(U_{1,g} \otimes U_{2,g})(v_1 \otimes v_2) = U_{1,g}(v_1) \otimes U_{2,g}(v_2).$$

We can use general statements about how operators extend to direct sums and tensor products to determine the characters of a composite representation L . Since $\chi_L(g) = \text{tr}(L_g)$, we need to determine the sum of the eigenvalues of L_g .

For a direct sum, the eigenvectors of $U_{1,g} \oplus U_{2,g}$ are $v_1 \oplus 0$, with eigenvalue λ_1 , and $0 \oplus v_2$, with eigenvalue λ_2 (where v_j is an eigenvector of $U_{j,g}$ with eigenvalue λ_j , etc.). So each eigenvalue of $U_{1,g}$ and $U_{2,g}$ contributes once. So, $\chi_{U_1 \oplus U_2}(g) = \text{tr}(U_{1,g}) + \text{tr}(U_{2,g}) = \chi_{U_1}(g) + \chi_{U_2}(g)$, i.e., $\chi_{U_1 \oplus U_2} = \chi_{U_1} + \chi_{U_2}$.

For a tensor product, the eigenvectors of $U_{1,g} \otimes U_{2,g}$ are $v_1 \otimes v_2$, with eigenvalue $\lambda_1 \lambda_2$. So every product of eigenvalues, one from V_1 and one from V_2 , contributes. So, $\chi_{U_1 \otimes U_2}(g) = \text{tr}(U_{1,g}) \text{tr}(U_{2,g}) = \chi_{U_1}(g) \chi_{U_2}(g)$, i.e., $\chi_{U_1 \otimes U_2} = \chi_{U_1} \chi_{U_2}$.

Breaking down representations

An “irreducible representation” V is one that cannot be broken down into a direct sum of two representations. One-dimensional representations, such as in the example for \mathbb{Z}_n , are necessarily irreducible.

Less obviously, for a commutative group – such as the translations of the line -- every irreducible representation is one-dimensional. The reason is the following. Let’s say you had some representation L of a commutative group that was of dimension 2 or more, and two group elements, say, g and h , for which the linear transformations L_g and L_h did *not* have a common set of eigenvectors. (If all L_x had a common set of eigenvectors, we could use this as a basis and decompose L into one-dimensional components.) If L_g and L_h did *not* have a common set of eigenvectors, then it would have to be that $L_g L_h \neq L_h L_g$, which would be a contradiction because $L_g L_h = L_{gh} = L_{hg} = L_h L_g$.

As a first step in breaking down a representation, we can ask whether there is any part of it that is trivial. That is, does V have a subspace, say W , for which L_g acts like the identity element? It turns out that we can find W by creating a projection P_L from V onto W . We define this projection as follows:

$$P_L(v) = \frac{1}{|G|} \sum_g L_g(v). \tag{7}$$

That is, we let every L_g act on a vector v , and average the result. Intuitively, the average vector $P_L(v)$ cannot be altered by any further group action, e.g., by some L_h . Formally,

$$P_L(L_h v) = \frac{1}{|G|} \sum_g L_g(L_h v) = \frac{1}{|G|} \sum_g L_{gh}(v) = \frac{1}{|G|} \sum_u L_u(v) = P_L(v),$$

where in the next-to-the last step we’ve replaced $u = gh$, and observed that letting g run over all of G is the same as letting $u = gh$ run over all of G .

The trace of a projection P is the dimension of the space that it projects onto. That is because,

when expressed in the basis of its eigenvalues, it looks like
$$\begin{pmatrix} 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & \cdots & 1 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 0 \end{pmatrix},$$
 where the 1’s

correspond to basis vectors that are unchanged by P (and span its range), and the 0’s correspond to the basis vectors that are set to 0.

So the dimension of W , the space on which L acts trivially, is given by $\text{tr}(P_L)$. This yields the “trace formula”:

$$\text{tr}(P_L(v)) = \frac{1}{|G|} \sum_g \text{tr}(L_g(v)) = \frac{1}{|G|} \sum_g \chi_L(g). \quad (8)$$

The regular representation

The “regular representation” is a representation that we are guaranteed to have for any group, and it arises from considering how the group acts on functions on a set, when the set itself is the group. To build the regular representation: Let V be the vector space of functions $x(g)$ from G to \mathbb{C} . (This is the “free vector space” on G). We can make V into a Hilbert space by defining $\langle x, y \rangle = \sum_G x(g) \overline{y(g)}$.

Note that this makes sense for infinite groups – our Hilbert space then consists of functions on the group for which the inner product of a function with itself is finite. For infinite but discrete groups (such as the integers, under addition), the above expression works fine. For infinite but continuous groups (such as the reals, under addition), we instead use $\langle x, y \rangle = \int_G x(g) \overline{y(g)} dg$.

We define the regular representation R as follows:

For each element p of G , we need to define R_p , a member of $\text{Hom}(V, V)$. R_p takes x (a function on G) to the $R_p(x)$ (another function on G) whose value at g is given by

$$(R_p(x))(g) = x(gp). \quad (9)$$

To see that R_p is unitary:

$\langle R_p(x), R_p(y) \rangle = \sum_{g \in G} (R_p(x))(g) \overline{(R_p(y))(g)} = \sum_{g \in G} x(gp) \overline{y(gp)} = \sum_{h \in G} x(h) \overline{y(h)}$. The reason for the final equality is that as g traverses G , then so does gp (but in a different order). Formally, change variables to $h = gp$; $g = hp^{-1}$ if hp^{-1} takes each value in G once, then so does h .

To see that R_p is a representation – i.e., that $R_p R_q = R_{pq}$: Here we are using the convention that $R_p R_q x$ means, “apply R_p to the result of applying R_q to x ”. So we need to show that

$R_p(R_q(x)) = R_{pq}(x)$ by evaluating the left and right hand side at every group element g .

On the left, say $y = R_q(x)$, so $y(g) = (R_q(x))(g) = x(gq)$. Then

$(R_p(y))(g) = y(gp) = x(gpq)$. On the right, $(R_{pq}(x))(g) = x(gpq)$.

Note that time translation as defined by (3) is an example of this: it is the regular representation of the additive group of the real numbers.

For a finite group, we can readily determine the character of the regular representation, as follows. We choose, as a basis for V , the functions on the group v_q , where $v_q(q) = 1$ and $v_q(g) = 0$ for $g \neq q$. R_p acts on V by permuting the v_q 's: $(R_p v_q)(g) = v_q(gp)$, which is nonzero only at $gp = q$, i.e., $g = qp^{-1}$. So $R_p v_q = v_{qp^{-1}}$. If $p = e$, the identity, then every v_q is mapped to itself, i.e., $\text{tr}(R_e) = \dim(V)$. But if $p \neq e$, every v_q is mapped to a different $v_{qp^{-1}}$. Viewed as a permutation matrix, R_p therefore must have its diagonal all 0's. So its trace is 0. Thus, for the regular representation, the character $\chi_R(s)$ is equal to 0 for all elements except the identity, and $\chi_R(e) = |G|$.

Finding parts of one representation inside another

With this last ingredient, we will see that the regular representation contains all irreducible representations of G .

The setup: a representation L of G in V (i.e., for each group element g , a unitary transformation L_g in $\text{Hom}(V, V)$), and another representation M of G in W (i.e., for each group element g , a unitary transformation M_g in $\text{Hom}(W, W)$). The statement that there is a part of L that corresponds to a part of M can be formalized by saying that there is a linear map φ in $\text{Hom}(V, W)$ for which $\varphi L_g = M_g \varphi$. That is, letting L act on a vector v in V , and then finding the image of $L_g(v)$ in W , is the same as finding the image $\varphi(v)$ of v in W , and letting M_g act on it.

So we want to know, whether any such φ 's exist.

The above setup allows us to define a group representation Φ in $\text{Hom}(V, W)$. The group representation is specified by Φ_g , which must map every element ϕ of $\text{Hom}(V, W)$ into some other element $\Phi_g(\phi)$. $\Phi_g(\phi)$ in turn must be specified by its action on any v . We choose:

$\Phi_g(\phi)(v) = M_g \phi L_g^{-1}(v)$. There are a few things to check – the most important of which is that it is a representation. That is, does $\Phi_{hg} = \Phi_h \Phi_g$? Making use of the fact that both M and L are group representations:

$$\begin{aligned} \Phi_{hg}(\phi)(v) &= M_{hg} \phi L_{hg}^{-1}(v) = M_{hg} \phi L_{(hg)^{-1}}(v) = M_{hg} \phi L_{g^{-1}h^{-1}}(v) = M_h M_g \phi L_{g^{-1}} L_{h^{-1}}(v) \\ &= M_h \Phi_g(\phi) L_{h^{-1}}(v) = \Phi_h(\Phi_g(\phi))(v) \end{aligned}$$

Notice that if every Φ_g acts as the identity on ϕ , then $M_g \phi L_g^{-1} = \Phi_g(\phi) = \phi$, and therefore that

$M_g \phi = \phi L_g$. Put another way, each operator in $\text{Hom}(V, W)$ for which Φ acts like the identity corresponds to a way of matching a component of L to a component of M . That is, the dimension of this space, which we will call $d(L, M)$, is the number of ways we can match components of L to component of M . Since this dimension is the size of the space in which Φ acts trivially, we can find it by applying the trace formula, eq. (8), to Φ .

$$d(L, M) = \frac{1}{|G|} \sum_g \chi_\Phi(g).$$

Just like we could evaluate the $\chi_{L \otimes M}$ from $\text{tr}(L_g \otimes M_g) = \text{tr}(L_g) \text{tr}(M_g)$, we can do the same for the representation Φ constructed in $\text{Hom}(V, W)$. (Or, one could use the correspondence between $\text{Hom}(V, W)$ and $V^* \otimes W$ -- see question 2D of Homework 3, 2008-2009 notes on Groups, Fields, and Vector Spaces -- and we can find a dual representation to L in $\text{Hom}(V^*, V^*)$ from the representation L in $\text{Hom}(V, V)$.) Either way leads to

$$\chi_\Phi(g) = \text{tr}(\Phi_g(\phi)) = \text{tr}(L_g^{-1}) \text{tr}(M_g) = \text{tr}(L_g^*) \text{tr}(M_g) = \overline{\text{tr}(L_g)} \text{tr}(M_g) = \overline{\chi_L(g)} \chi_M(g).$$

This yields our main result:

$$d(L, M) = \frac{1}{|G|} \sum_g \overline{\chi_L(g)} \chi_M(g). \tag{10}$$

As a special case for $L = M$:

$$d(L, L) = \frac{1}{|G|} \sum_g |\chi_L(g)|^2. \tag{11}$$

The group representation theorem

While we have carried this analysis out for finite groups, everything we've done leading up to eq. (10) also works for infinite groups, provided that we can set up a Hilbert space in the functions on them (which amounts to being able to define integrals, so that there is a dot-product).

By applying eq. (10) to a few special cases, we obtain all the main properties of group representations, which are summarized in the "group representation theorem" -- and which formalizes the non-accidental nature of the Fourier transform.

Recall that an "irreducible representation" is a representation cannot be written as a direct sum of group representations.

Here are the facts:

- The characters of irreducible representations are orthogonal. This follows from eq. (10) directly, since (according to the definition of irreducible representations), if L and M are two different irreducible representations, $d(L, M) = 0$.
- The character of an irreducible representation is an orthonormal function on the group. This follows from eq. (11), since in this case, $d(L, L) = 1$.
- Every irreducible representation L occurs in the regular representation, and the number of occurrences is equal to the dimension of L . This follows from eq. (10) by taking M to be the regular representation, R . The character of the regular representation is 0 for all group elements except the identity, and is $|G|$ at the identity. So the only term that contributes to the sum is the term for $g = e$. $\chi_L(e)$ is the dimension of L , since the representation of e is the identity matrix (and the trace just adds up the 1's on the diagonal).

For finite, commutative groups, we can go further very easily by counting dimensions. Since every irreducible representation is one-dimensional, the number of different irreducible representations must be $|G|$. Thus, the characters of the irreducible representations form an orthonormal basis for functions on the group.

Algebraically, nothing changes when one goes from finite groups to infinite ones, but there are things to prove (about limits, integrals, etc.)

Ignoring these “details”, we apply the above to the additive group of the real numbers. Its regular representation is the time translation operators defined by eq. (3). All irreducible representations must be one-dimensional. Above we showed that each representation must be of the form $T \rightarrow e^{i\omega T}$. So this is the full set, and we have decomposed space of the regular representation (the space of functions of time) into one-dimensional subspaces, in which time translation by T acts like multiplication by $e^{i\omega T}$,

For finite but non-commutative groups, it is a bit more complex. There will always be some conjugate classes with more than one element, since there will be always some choice of g and h for which $hgh^{-1} \neq g$. So there have to be conjugate classes than $|G|$ (since at least one of them has two or more elements). So there have to be fewer distinct irreducible representations than $|G|$, since their characters must be orthogonal (and hence, linearly independent) functions on the conjugate classes. With a bit more work, one can show that the matrix elements of the irreducible representations are orthonormal functions on the group (look at the group-average of the tensor product of two representations).

Example: the representations of the cyclic group

To get an idea of what happens in the commutative case, here we consider a generic cyclic group \mathbb{Z}_n . We can regard this as the group generated by a single element a , of order n .

Since it is commutative, then all irreducible representations are one-dimensional. A unitary 1×1 matrix is simply a complex number of magnitude 1. Say a maps to the complex number z .

Since $a^n = e$, it follows that $z^n = 1$, i.e., that $z = \exp\left(\frac{2\pi i}{n} m\right)$ for some m . Each choice of m in $\{0, 1, \dots, n-1\}$ yields a different group representation, as it yields a distinct z . Since there are n such choices, we have found all the irreducible representations.

Summing up: the m th representation L_m is: $L_m(a) = \exp\left(\frac{2\pi i}{n} m\right)$ (and $L_m(a^j) = \exp\left(\frac{2\pi i}{n} mj\right)$),

and its character is $\chi_{L_m}(a^k) = \exp\left(\frac{2\pi i}{n} mj\right)$. Writing a function on the group elements as a sum of the characters is the discrete Fourier transform.

The orthonormality guaranteed by the group representation theory is that $d(L_m, L_p) = 0$ for $m \neq p$ and $d(L_m, L_m) = 1$, where

$d(L_m, L_p) = \frac{1}{n} \sum_{j=0}^{n-1} \overline{\chi_m(a^j)} \chi_p(a^j)$. This can be seen directly:

$d(L_m, L_p) = \frac{1}{n} \sum_{j=0}^{n-1} \exp\left(-\frac{2\pi i}{n} mj\right) \exp\left(\frac{2\pi i}{n} pj\right) = \frac{1}{n} \sum_{j=0}^{n-1} \exp\left(\frac{2\pi i}{n} (p-m)j\right)$. If $p = m$, all terms on the right hand side are 1. If $p \neq m$, the right side is a symmetric sum over distinct roots of unity.

Example: the representations of the group of the cube

To get an idea of what happens in the non-commutative case, here we consider the group of the rotations of a standard 3-d cube. Since we can move any of its six faces into a standard position, and then rotate in any of 4 steps, this group has 24 elements. (Abstractly, this is also the same as the group of permutations of 4 elements – but we won't use that fact explicitly – you can see this by thinking about how rotations of the cube act on its four diagonals.)

Here we work out its “character table” – i.e., a table of the characters of all of its representations. It illustrates many of the properties of characters and representations.

The first step is determining the conjugate classes – these are the sets containing group elements that are identical up to inner automorphism. I.e., if two group elements are the same except for a relabeling due to rotation of the cube, they are in the same conjugate class.

We label three cube faces A, B, C , and their opposites $A', B',$ and C' .

- (1) The identity – always one element in this class.
- (2) 90-deg rotations around a face – 6 faces, so 6 elements. This is the same as considering 90-deg clockwise or counterclockwise rotations around the axes AA' , BB' , or CC' .
- (3) 180-deg rotations around an axis – 3 elements
- (4) 120-deg rotations around a vertex – 8 elements
- (5) 180-deg rotations around the midpoint of opposite edges – e.g., exchange A with B , A' with B' , and C with C' . – 6 elements (there are 12 edges, so 6 pairs of edges to do this with)

As a check, we now have all 24 elements (1+6+3+8+6), in 5 conjugate classes.

Building the character table

We begin to write the “character table” by setting up a header row with the conjugate classes, and subsequent rows to contain the characters. The numbers in square brackets indicate the number of elements in the conjugate class. The first row is the trivial representation; it is one-dimensional and, since it maps each group element into 1, its character is 1.

	<u>$id[1]$</u>	<u>$face90[6]$</u>	<u>$face180[3]$</u>	<u>$vertex120[8]$</u>	<u>$edge180[6]$</u>
E	1	1	1	1	1

To find some other representations:

Every group element permutes the face-pairs – AA' , BB' , or CC' . They can thus be represented as permutation matrices on the three items, AA' , BB' , or CC' . Let’s call this representation F (for faces). $\chi_F(g)$, which is the number of elements on the diagonal of F_g , is the number of face-pairs that are unchanged by the group element. For the identity, all are unchanged, so the character is 3. For a 90-deg rotation, one face-pair is unchanged and the other two are swapped, so the character is 1. For a 180-degree rotation, they’re all preserved, so the character is 3. For the 120-deg rotation, they are cycled, so the character is 0 (none are preserved). For the edge-flip (e.g., around an edge between an A -face and a B -face), two pairs are interchanged, and the third pair is preserved, so the character is 1.

	<u>$id[1]$</u>	<u>$face90[6]$</u>	<u>$face180[3]$</u>	<u>$vertex120[8]$</u>	<u>$edge180[6]$</u>
F	3	1	3	0	1

We now need to check whether F is irreducible. According to the group representation theorem, it is irreducible if $d(F, F) = 1$. So we calculate (using eq. (10)), and using the numbers in the square brackets to keep track of the number of elements in each conjugate class:

$$d(F, F) = \frac{1}{24}(1 \cdot 3^2 + 6 \cdot 1^2 + 3 \cdot 3^2 + 8 \cdot 0^2 + 6 \cdot 1^2) = \frac{48}{24} = 2, \text{ i.e., } F \text{ is not irreducible.}$$

Perhaps F contains a copy of E . If so, we can remove that copy (by $I - P_F$, where P_F is given by eq. (7)), and find a new irreducible representation. To see if F contains a copy of E , we calculate (using eq. (10)):

$$d(E, F) = \frac{1}{24}(1 \cdot (1 \cdot 3) + 6 \cdot (1 \cdot 1) + 3 \cdot (1 \cdot 3) + 8 \cdot (1 \cdot 0) + 6 \cdot (1 \cdot 1)) = \frac{24}{24} = 1.$$

This means that F contains E . (It wasn't a lucky guess; since the character of F was non-negative, then $d(E, F)$ had to be > 0 .) To find the other part of F , we could work out $I - P_F$ (using eq. (7)), to project onto a subspace that contains no copies of E – and hence, which contains some other representation, say F_0 , with $F = E \oplus F_0$. But it is easier just to compute the character of F_0 : $\chi_F = \chi_E + \chi_{F_0}$, so $\chi_{F_0} = \chi_F - \chi_E$. Entering this into the table:

	<u>id[1]</u>	<u>face90[6]</u>	<u>face180[3]</u>	<u>vertex120[8]</u>	<u>edge180[6]</u>
E	1	1	1	1	1
F_0	2	0	2	-1	0

Another representation is simply regarding these group elements as 3-d rotations, and writing them as 3×3 matrices. Let's call this M . To determine the character, we only need to write the matrix out for one example of each conjugate class, since the character is constant on conjugate classes. The identity group element of course yields the identity 3×3 matrix, and a character of

3. A 90-deg face rotation that rotates in the XY -plane has a matrix $\begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, and a character

of 1. A 180-deg rotation, which is the square of this matrix, has matrix $\begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, and a

character of -1. A 120-deg rotation around a vertex permutes the axes, and so has matrix

$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$, and character 0. An edge-flip could exchange X for Y , and Y for X , and invert

Z, and thus, have matrix $\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}$, and character -1. Using eq. (11), we find that M is

irreducible:

$$d(M, M) = \frac{1}{24} (1 \cdot 3^2 + 6 \cdot 1^2 + 3 \cdot (-1)^2 + 8 \cdot 0^2 + 6 \cdot (-1)^2) = \frac{24}{24} = 1. \text{ Adding this to the table:}$$

	<u>id[1]</u>	<u>face90[6]</u>	<u>face180[3]</u>	<u>vertex120[8]</u>	<u>edge180[6]</u>
E	1	1	1	1	1
F_0	2	0	2	-1	0
M	3	1	-1	0	-1

As noted above, every group element acts on the three sets of axes, and permutes them. Some group elements cause an odd permutation of the axes (i.e., swap one pair), while others lead to an even permutation (i.e., don't swap any axes, or, cycle through all three of them). So there is a group representation Q that maps each group element to -1 or 1, depending on whether the permutation of the axes is odd or even. Since this is a one-dimensional representation, it must be irreducible. Adding it to the table:

	<u>id[1]</u>	<u>face90[6]</u>	<u>face180[3]</u>	<u>vertex120[8]</u>	<u>edge180[6]</u>
E	1	1	1	1	1
F_0	2	0	2	-1	0
M	3	1	-1	0	-1
Q	1	-1	1	1	-1

Now let's use tensor products to create a representation. $M \otimes Q$ is a good choice: Q is all 1's and -1's, so if M is irreducible, then so will $M \otimes Q$. (See eq. (11): $\chi_{M \otimes Q} = \chi_M \chi_Q$, so

$$|\chi_{M \otimes Q}|^2 = |\chi_M|^2, \text{ so } d(M \otimes Q, M \otimes Q) = d(M, M) = 1.) \text{ Adding this to the table:}$$

	<u>id[1]</u>	<u>face90[6]</u>	<u>face180[3]</u>	<u>vertex120[8]</u>	<u>edge180[6]</u>
E	1	1	1	1	1
F_0	2	0	2	-1	0
M	3	1	-1	0	-1
Q	1	-1	1	1	-1
$M \otimes Q$	3	-1	-1	0	1

The table is now finished. We can verify that we've fully decomposed the regular representation – it should have each irreducible representation, repeated a number of times equal to the dimension of the representation, and indeed, $24 = 1^2 + 2^2 + 3^2 + 1^2 + 3^2$. One can also verify that the rows are orthogonal (and columns too!).

If you try to make new representations by tensoring these, you don't get anything new. For example (verify using the characters), $F_0 \otimes F_0 = E \oplus Q \oplus F_0$.

Character tables can contain non-integer values, as is the case for \mathbb{Z}_n .