

Recapitulating main definitions and results

For a linear transformation A in a vector space V , an eigenvector 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 .

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

A linear transformation is typically associated with a *set* of eigenvalues λ_j and their associated eigenvectors v_j , satisfying

$$Av_j = \lambda_j v_j. \quad (1)$$

For a finite-dimensional vector space V , we can find the eigenvalues of A by solving the equation

$$\det(A - zI) = 0 \quad (2)$$

For a vector space V of dimension n , eq. (2) is a polynomial of degree n . So if the field k is algebraically closed, solutions of eq. (2) will exist. This is why we choose $k = \mathbb{C}$, the field of complex numbers: \mathbb{C} is algebraically closed.

The significance of this is that the eigenvalues and eigenvectors 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 lie in a common subspace.

If 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 . To see this, note $A(av + bw) = aAv + bAw = a\lambda v + b\lambda w = \lambda(av + bw)$. 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 whose eigenvalues are $\lambda_m \neq \lambda$.

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, $Av = \lambda v$ 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.

When the eigenvectors form a basis

Say there is a special linear transformation T (e.g., one specified by the problem at hand) and its eigenvectors v_j (with eigenvalues λ_j) formed a basis. Then, the action of T on any vector $v \in V$ could be specified: since $v = \sum a_j v_j$ for some set of coefficients a_j , then $T(v) = T(\sum a_j v_j) = \sum T(a_j v_j) = \sum a_j T(v_j) = \sum a_j \lambda_j v_j$. Another way of looking at this is that if you use the eigenvectors v_j as the basis set, then the matrix

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

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 , and, we can even interpret $f(T)$ as a transformation with eigenvalues $f(\lambda_j)$, for any function f .

Note that there is no guarantee that the eigenvalues of a linear transformation form a basis. See Homework Q1. But there is a guarantee that eigenvectors corresponding to distinct eigenvalues are linearly independent.

Shared eigenvectors and commutation

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

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

If, furthermore, the eigenspace of B corresponding to eigenvalue λ has 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 .

Our setup

V 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

$$w(t) = \int_0^{\infty} L(\tau)v(t-\tau)d\tau \quad (3)$$

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

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

Example: smoothing

$w = Lv$ is a smoothing transformation for

$$w(t) = \int_{-\infty}^{\infty} L(\tau)v(t-\tau)d\tau \quad (4)$$

and $L(t) = 1/(2h)$ if $|t| < h$, 0 otherwise. L in this context is often called the “smoothing kernel.”

Other examples will arise when we discuss point processes.

Time-translation invariance

The above examples are “time-translation invariant.” That is, they are independent of absolute clock time. This is a crucial property. It can be formulated algebraically. First, we define the time-shift operator (linear transformation) D_T as follows:

$$(D_T v)(t) = v(t+T). \quad (5)$$

This is equivalent to an expression of the form (4), with $L(\tau) = \delta(\tau+T)$. In particular,

$$\int_{-\infty}^{\infty} \delta(\tau+T)v(t-\tau)d\tau = v(t+T) = (D_T v)(t),$$

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.

Since $AD_T = D_TA$, we can use the relationship of commuting operators and eigenvectors. That is, if we can find a basis set consisting of the eigenvectors of D_T , for all T , then we have found the eigenvectors for all time-translation invariant linear operators A , that is, a basis set in which all linear transformations of the form (3) or (4) are diagonal.

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 $c = i\omega$, $v(T) = e^{i\omega T}$.

Based on some very general results, the set of such $v(T) = e^{i\omega T}$ (for all ω) form the complete set of eigenvectors of each of the D_T 's, and 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. (3)).

Hilbert spaces

To see why this happens, we need one more piece of structure to vector spaces: the inner product. An inner product, essentially, adds the notion of distance. A vector space with an inner product 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.

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 $v \bullet w$. 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 $k = \mathbb{C}$.

Linearity: $\langle av_1 + bv_2, w \rangle = a\langle v_1, w \rangle + b\langle v_2, w \rangle$, and $\langle v, aw_1 + bw_2 \rangle = \overline{a}\langle v, w_1 \rangle + \overline{b}\langle v, w_2 \rangle$ (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.

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

Note that $\langle av, bw \rangle = a\overline{b}\langle v, w \rangle$. The necessity for the complex-conjugation is apparent if one considers $\langle iv, iv \rangle$. With complex-conjugation of the “ b ”, we find

$\langle iv, iv \rangle = i\overline{i}\langle v, v \rangle = i(-i)\langle v, v \rangle = \langle v, v \rangle$, 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.

The quantity specified by $d(v, w) = \|v - w\| = \sqrt{\langle v - w, v - w \rangle}$ satisfies the triangle inequality $d(u, w) \leq d(u, v) + d(v, w)$, and it is symmetric and non-negative – and therefore qualifies as a “metric” (i.e., a distance).

For the vector space V of functions of time, the standard inner product is

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(t)\overline{g(t)}dt \quad (6)$$

and the corresponding Hilbert space is the subset of V for which the integral for $\langle f, f \rangle$ exists and is finite.

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

$$\langle u, v \rangle = \sum_{n=1}^N u_n \overline{v_n}. \quad (7)$$

Note that although we used coordinates to define these inner product, defining an inner product is not the same as specifying coordinates. The inner product only fixes a notion of distance, while the coordinates specify individual directions.

An inner product is equivalent to specifying a correspondence between a vector space V and its dual V^* . 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 $\varphi_{av} = \overline{a}\varphi_v$.

Adjoins

The adjoint of an operator A is the operator A^* (sometimes written A^\dagger) for which $\langle Au, v \rangle = \langle u, A^*v \rangle$.

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

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

Some special kinds of linear operators

A **self-adjoint** operator A is an operator for which $A = A^*$. Self-adjoint operators have real eigenvalues. This follows from noting that if $Av = \lambda v$, then

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

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

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

$$\lambda \langle v, w \rangle = \langle \lambda v, w \rangle = \langle Av, w \rangle = \langle v, A^*w \rangle = \langle v, Aw \rangle = \langle v, \mu w \rangle = \overline{\mu} \langle v, w \rangle, \text{ so } \lambda = \overline{\mu} \text{ or}$$

$\langle v, w \rangle = 0$. Since both λ and μ are real, and they are assumed to be unequal, it follows that $\langle v, w \rangle = 0$.

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. This 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

$$\langle v, w \rangle = \langle Av, Aw \rangle = \langle \lambda v, \mu w \rangle = \lambda \bar{\mu} \langle v, w \rangle = \frac{\lambda}{\mu} \langle v, w \rangle \text{ (with the last equality because } \mu \bar{\mu} = |\mu|^2 = 1 \text{). So if } \lambda \neq \mu, \text{ then } \langle v, w \rangle = 0.$$

Note that the time-translation operator D_T is unitary.

A **projection** operator is a self-adjoint operator P for which $P^2 = P$. One should 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 $Q^2 = (I - P)^2 = (I - P)(I - P) = I - IP - PI + P^2 = I - P - P + P = I - P = Q$. Also $PQ = P(I - P) = P - P^2 = 0$. 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 $P_u(v) = v - u \frac{\langle v, u \rangle}{\langle u, u \rangle}$.

To see that P_u this is self-adjoint, note that

$$\langle P_u(v), w \rangle = \langle v, w \rangle - \langle u, w \rangle \frac{\langle v, u \rangle}{\langle u, u \rangle} = \langle v, w \rangle - \frac{\langle v, u \rangle \langle w, u \rangle}{\langle u, u \rangle} \text{ but also}$$

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

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

$$P_u^2(v) = P_u(v) - u \frac{\langle P_u(v), u \rangle}{\langle u, u \rangle} = v - u \frac{\langle v, u \rangle}{\langle u, u \rangle} - \frac{u}{\langle u, u \rangle} \left\langle v - u \frac{\langle v, u \rangle}{\langle u, u \rangle}, u \right\rangle. \text{ The last term drops}$$

out because

$$\left\langle v - u \frac{\langle v, u \rangle}{\langle u, u \rangle}, u \right\rangle = \langle v, u \rangle - \langle u, u \rangle \frac{\langle v, u \rangle}{\langle u, u \rangle} = 0.$$

A **normal** operator is an operator that commutes with its adjoint. Self-adjoint and unitary operators are normal.

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

$$A = \sum_{\lambda} \lambda P_{\lambda} \quad (8)$$

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. (8)) 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.”

Unitary representations: definition and simple example

A unitary representation U of a group G is a group isomorphism from elements g of G into unitary operators U_g in $\text{Hom}(V, V)$. 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 $\text{Hom}(V, V)$.

A simple example: 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

$$U_p U_q = e^{\frac{2\pi i}{n}p} e^{\frac{2\pi i}{n}q} = e^{\frac{2\pi i}{n}(p+q)} = U_{p+q}.$$

Two representations, U_1 in V_1 and U_2 in V_2 , can be combined to make a composite representation in $V_1 \oplus V_2$. An “irreducible representation” V is one that cannot be broken down into a such a direct sum of two representations. One-dimensional representations, such as the example above, are necessarily irreducible.

Less obviously, for a commutative group, every irreducible representation is one-dimensional. To see this (plausibility argument, not a proof), we note that if there is a non-trivial operator A that commutes with all of the group representation operators U_g , then U can be reduced – namely, into the eigenspaces of A .

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 $\langle x, y \rangle = \sum_G x(g) \overline{y(g)}$. If G is

continuous, we instead use $\langle x, y \rangle = \int_G x(g) \overline{y(g)} dg$. We define the regular representation

R as follows:

R_p , a member of $\text{Hom}(V, V)$, 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 (5) is an example of this: it is the regular representation of the additive group of the real numbers.

The group representation theorem

The regular representation contains at least one copy of every irreducible representation. The number of copies of an irreducible representation U is equal to the dimension of U .

We apply this to the additive group of the real numbers. Its regular representation is the time translation operators defined by (5). 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 $e^{i\omega T}$,