

Linear Algebra for Quantum Mechanics

Introduction

In quantum mechanics, the state of an electron in some potential is given by a wave function $\psi(\vec{x}, t)$, and physical variables are represented by operators on this wave function, such as momentum in the x-direction $p_x = -i\hbar\partial/\partial x$.

The Schrodinger wave equation is a *linear* equation, which means that if ψ_1 and ψ_2 are solutions, then so is $c_1\psi_1 + c_2\psi_2$, where c_1, c_2 are arbitrary complex numbers.

This linearity of the sets of possible solutions is true generally in quantum mechanics, as is the representation of physical variables by operators on the wave functions.

The mathematical structure this describes, the linear set of possible states and sets of operators on those states, is in fact a *linear algebra* of operators acting on a *vector space*.

From now on, this is the language we'll be using most of the time.

To clarify, we'll give some definitions.

What is a Vector Space?

The prototypical vector space is of course the set of real vectors in ordinary three-dimensional space, these vectors can be represented by trios of real numbers (v_1, v_2, v_3) measuring the components in the x, y and z directions respectively.

The basic properties of these vectors are:

- any vector multiplied by a number is another vector in the space,

$$a(v_1, v_2, v_3) = (av_1, av_2, av_3)$$

- the sum of two vectors is another vector in the space, that given by just adding the corresponding components together: $(v_1 + w_1, v_2 + w_2, v_3 + w_3)$

These two properties together are referred to as “***closure***”: adding vectors and multiplying them by numbers cannot get you out of the space.

- A further property is that there is a unique null vector $(0, 0, 0)$ and each vector has an additive inverse $(-v_1, -v_2, -v_3)$ which added to the original vector gives the null vector.

Mathematicians have generalized the definition of a vector space: a general vector space has the properties we’ve listed above for three-dimensional real vectors, but the operations of addition and multiplication by a number are generalized to more abstract operations between more general entities.

The operators are, however, restricted to being commutative and associative.

Notice that the list of necessary properties for a general vector space does not include that the vectors have a magnitude—that would be an additional requirement, giving what is called a *normed* vector space.

More about that later.

To go from the familiar three-dimensional vector space to the vector spaces relevant to quantum mechanics, first the real numbers (components of the vector and possible multiplying factors) are to be generalized to complex numbers, and second the three-component vector goes to an n component vector.

The consequent n -dimensional complex space is sufficient to describe the quantum mechanics of angular momentum, an important subject.

But to describe the wave function of a particle in a box requires an *infinite* dimensional space, one dimension for each Fourier component, and to describe the wave function for a particle on an infinite line requires the set of all normalizable continuous differentiable functions on that line.

Fortunately, all these generalizations are to finite or infinite sets of complex numbers, so the mathematicians' vector space requirements of commutativity and associativity are always trivially satisfied.

We use Dirac's notation for vectors, $|1\rangle$, $|2\rangle$ and call them "kets", so, in his language, if $|1\rangle$, $|2\rangle$ belong to the space, so does $c_1|1\rangle + c_2|2\rangle$ for arbitrary complex constants c_1, c_2 .

Since our vectors are made up of complex numbers, multiplying any vector by zero gives the null vector, and the additive inverse is given by reversing the signs of all the numbers in the vector.

The set of solutions of Schrödinger's equation for an electron in a potential satisfies the requirements for a vector space: $\psi(\vec{x}, t)$ is just a complex number at each point in space, so only complex numbers are involved in forming $c_1\psi_1 + c_2\psi_2$, and commutativity, associativity, etc., follow at once.

Vector Space Dimensionality

The vectors $|1\rangle, |2\rangle, |3\rangle$ are *linearly independent* if

$$c_1 |1\rangle + c_2 |2\rangle + c_3 |3\rangle = 0$$

implies

$$c_1 = c_2 = c_3 = 0$$

A vector space is n -dimensional if the maximum number of linearly independent vectors in the space is n .

Such a space is often called $V^n(C)$ or $V^n(R)$ if only real numbers are used.

Now, vector spaces with finite dimension n are clearly insufficient for describing functions of a continuous variable x .

But they are well worth reviewing here: as we've mentioned, they are fine for describing quantized angular momentum, and they serve as a natural introduction to the infinite-dimensional spaces needed to describe spatial wavefunctions.

A set of n linearly independent vectors in n -dimensional space is a *basis*: any vector can be written *in a unique way* as a sum over a basis:

$$|V\rangle = \sum_i v_i |i\rangle$$

You can check the uniqueness by taking the difference between two supposedly distinct sums: it will be a linear relation between independent vectors, a contradiction.

Since all vectors in the space can be written as linear sums over the elements of the basis, the sum of multiples of any two vectors has the form:

$$a|V\rangle + b|W\rangle = \sum_i (av_i + bw_i) |i\rangle$$

Inner Product Spaces

The vector spaces of relevance in quantum mechanics also have an operation associating a number with a pair of vectors, a generalization of the dot product of two ordinary three-dimensional vectors,

$$\vec{a} \cdot \vec{b} = \sum_i a_i b_i$$

Following Dirac, we write the inner product of two ket vectors $|V\rangle, |W\rangle$ as $\langle W | V \rangle$. Dirac refers to this $\langle | \rangle$ form as a “bracket” made up of a “bra” and a “ket”.

This means that each ket vector $|V\rangle$ has an associated bra $\langle V|$.

For the case of a real n -dimensional vector, $|V\rangle, \langle V|$ have identical components - but we require for the more general case that

$$\langle W | V \rangle = \langle V | W \rangle^*$$

where $*$ denotes complex conjugate.

This implies that for a ket $(v_1, \dots, v_n)$ the bra will be $(v_1^*, \dots, v_n^*)$.

Actually, bras are usually written as rows, kets as columns, so that the inner product follows standard rules for matrix multiplication.

Evidently for the n -dimensional complex vector $\langle V | V \rangle$ is real and positive except for the null vector:

$$\langle V | V \rangle = \sum_{i=1}^n |v_i|^2$$

For the more general inner product spaces considered later we require $\langle V | V \rangle$ to be positive, except for the null vector. These requirements do restrict the classes of vector spaces we are considering - no Lorentz metric, for example - but they are all satisfied by the spaces relevant to nonrelativistic quantum mechanics.)

The *norm* of $|V\rangle$ is then defined by

$$|V| = \sqrt{\langle V | V \rangle}$$

If $|V\rangle$ is a member of $V^n(C)$, so is $a|V\rangle$, for any complex number a .

We require the inner product operation to commute with multiplication by a number, so

$$\langle W | (a|V\rangle) = a \langle W | V \rangle$$

The complex conjugate of the right hand side is $a^* \langle V | W \rangle$.

For consistency, the bra corresponding to the ket $a|V\rangle$ must therefore be $\langle V | a^*$ - in any case obvious from the definition of the bra in n complex dimensions given above.

It follows that if

$$|V\rangle = \sum_i v_i |i\rangle, \quad |W\rangle = \sum_i w_i |i\rangle, \quad \text{then} \quad \langle V | W \rangle = \sum_{i,j} v_i^* w_j \langle i | j \rangle$$

Constructing an Orthonormal Basis: the Gram-Schmidt Process

To have something better resembling the standard dot product of ordinary three vectors, we need $\langle i | j \rangle = \delta_{ij}$, that is, we need to construct an *orthonormal basis* in the space.

There is a straightforward procedure for doing this called the Gram-Schmidt process.

We begin with a linearly independent set of basis vectors, $|1\rangle, |2\rangle, |3\rangle, \dots$.

We first normalize $|1\rangle$ by dividing it by its norm. Call the normalized vector $|I\rangle$.

Now $|2\rangle$ cannot be parallel to $|I\rangle$, because the original basis was of linearly independent vectors, but $|2\rangle$ in general has a nonzero component parallel to $|I\rangle$, equal to $|I\rangle \langle I | 2 \rangle$, since $|I\rangle$ is normalized.

Therefore, the vector $|2\rangle - |I\rangle \langle I | 2 \rangle$ is perpendicular to $|I\rangle$, as is easily verified.

It is also easy to compute the norm of this vector, and divide by it to get $|II\rangle$, the second member of the orthonormal basis.

Next, we take $|3\rangle$ and subtract off its components in the directions $|I\rangle$ and $|II\rangle$, normalize the remainder, and so on.

In an n -dimensional space, having constructed an orthonormal basis with members $|i\rangle$, any vector $|V\rangle$ can be written as a column vector,

$$|V\rangle = \sum_i v_i |i\rangle = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ \vdots \\ v_n \end{pmatrix} \quad \text{where} \quad |1\rangle = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ \vdots \\ 0 \end{pmatrix} \quad \text{and so on.}$$

The corresponding bra is $\langle V| = \sum v_i^* \langle i|$, which we write as a row vector with the elements complex conjugated, $\langle V| = (v_1^*, v_2^*, \dots, v_n^*)$.

This operation, going from columns to rows and taking the complex conjugate, is called taking the *adjoint*, and can also be applied to matrices, as we shall see shortly.

The reason for representing the bra as a row is that the inner product of two vectors is then given by standard matrix multiplication:

$$\langle V | W \rangle = (v_1^*, v_2^*, \dots, v_n^*) \begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ \vdots \\ w_n \end{pmatrix}$$

Of course, this only works with an orthonormal base.

The Schwartz Inequality

The Schwartz inequality is the generalization to any inner product space of the result

$$|\vec{a} \cdot \vec{b}|^2 \leq |\vec{a}|^2 |\vec{b}|^2 \quad \text{or} \quad \cos^2 \theta \leq 1$$

for ordinary three-dimensional vectors.

The equality sign in that result only holds when the vectors are parallel.

To generalize to higher dimensions, one might just note that two vectors are in a two-dimensional subspace, but an illuminating way of understanding the inequality is to write the vector $\vec{a}$ as a sum of two components, one parallel to $\vec{b}$ and one perpendicular to $\vec{b}$.

The component parallel to $\vec{b}$ is just $\vec{b}(\vec{a} \cdot \vec{b})/|\vec{b}|^2$, so the component perpendicular to $\vec{b}$ is the vector $\vec{a}_\perp = \vec{a} - \vec{b}(\vec{a} \cdot \vec{b})/|\vec{b}|^2$.

Substituting this expression into $\vec{a}_\perp \cdot \vec{a}_\perp \geq 0$, the equality follows.

This same point can be made in a general inner product space: if $|V\rangle, |W\rangle$ are two vectors, then

$$|Z\rangle = |V\rangle - \frac{|W\rangle \langle W | V \rangle}{|w|^2}$$

is the component of $|V\rangle$ perpendicular to $|W\rangle$, as is easily checked by taking its inner product with $|W\rangle$.

Then

$$\langle Z | Z \rangle \geq 0 \quad \text{gives immediately} \quad |\langle V | W \rangle|^2 \leq |V|^2 |W|^2$$

Linear Operators

A *linear operator* A takes any vector in a linear vector space to a vector in that space, $A|V\rangle = |V'\rangle$, and satisfies

$$A(c_1|V_1\rangle + c_2|V_2\rangle) = c_1A|V_1\rangle + c_2A|V_2\rangle$$

with c_1, c_2 arbitrary complex constants.

The *identity operator* I is (obviously!) defined by:

$$I|V\rangle = |V\rangle \quad \text{for all } |V\rangle$$

For an n -dimensional vector space with an orthonormal basis $|1\rangle, \dots, |n\rangle$, since any vector in the space can be expressed as a sum $|V\rangle = \sum v_i |i\rangle$, the linear operator is completely determined by its action on the basis vectors - this is all we need to know.

It's easy to find an expression for the identity operator in terms of bras and kets.

Taking the inner product of both sides of the equation $|V\rangle = \sum v_i |i\rangle$ with the bra $\langle i|$ gives $\langle i|V\rangle = v_i$, so

$$|V\rangle = \sum_i v_i |i\rangle = \sum_i |i\rangle \langle i|V\rangle$$

Since this is true for any vector in the space, it follows that that the identity operator is just

$$I = \sum_{i=1}^n |i\rangle \langle i|$$

This is an important result: it will reappear in many disguises.

To analyze the action of a general linear operator A , we just need to know how it acts on each basis vector. Beginning with $A|1\rangle$ this must be some sum over the basis vectors, and since they are orthonormal, the component in the $|i\rangle$ direction must be just $\langle i|A|1\rangle$.

That is,

$$A|1\rangle = \sum_{i=1}^n |i\rangle \langle i|A|1\rangle = \sum_{i=1}^n A_{i1} |i\rangle \quad , \text{ writing } \quad \langle i|A|1\rangle = A_{i1}$$

So if the linear operator A acting on $|V\rangle = \sum_i v_i |i\rangle$ gives $|V'\rangle = \sum_i v'_i |i\rangle$, that is $A|V\rangle = |V'\rangle$, the linearity tells us that

$$\sum_i v'_i |i\rangle = |V'\rangle = A|V\rangle = \sum_j v_j |j\rangle = \sum_{i,j} v_j |i\rangle \langle i|A|j\rangle = \sum_{i,j} v_j A_{ij} |i\rangle$$

where in the fourth step we just inserted the identity operator.

Since the $|i\rangle$'s are all orthogonal, the coefficient of a particular $|i\rangle$ on the left-hand side of the equation must be identical with the coefficient of the same $|i\rangle$ on the right-hand side. That is,

$$v'_i = A_{ij} v_j$$

Therefore the operator A is simply equivalent to *matrix multiplication*:

$$\begin{pmatrix} v'_1 \\ v'_2 \\ \vdots \\ \vdots \\ v'_n \end{pmatrix} = \begin{pmatrix} \langle 1| A |1\rangle & \langle 1| A |2\rangle & \cdots & \langle 1| A |n\rangle \\ \langle 2| A |1\rangle & \langle 2| A |2\rangle & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ \langle n| A |1\rangle & \langle n| A |2\rangle & \cdots & \langle n| A |n\rangle \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ \vdots \\ v_n \end{pmatrix}$$

Evidently, then, applying two linear operators one after the other is equivalent to successive matrix multiplication - and, therefore, *since matrices do not in general commute, nor do linear operators*.

Of course, if we hope to represent quantum variables as linear operators on a vector space, this has to be true: the momentum operator $p_x = -i\hbar d/dx$ certainly doesn't commute with x !)

Projection Operators

It is important to note that a linear operator applied successively to the members of an orthonormal basis might give a new set of vectors which *no longer span the entire space*.

To give an example, the linear operator $|1\rangle \langle 1|$ applied to any vector in the space picks out the vector's component in the $|1\rangle$ direction.

It's called a *projection operator*.

The operator $|1\rangle \langle 1| + |2\rangle \langle 2|$ projects a vector into its components in the subspace spanned by the vectors $|1\rangle$ and $|2\rangle$, and so on: if we extend the sum to be over the whole basis, we recover the identity operator.

Exercise: prove that the $n \times n$ matrix representation of the projection operator $|1\rangle\langle 1| + |2\rangle\langle 2|$ has all elements zero except the first two diagonal elements, which are equal to one.

There can be no *inverse* operator to a nontrivial projection operator, since the information about components of the vector perpendicular to the projected subspace is lost.

The Adjoint Operator and Hermitian Matrices

As we've discussed, if a ket $|V\rangle$ in the n -dimensional space is written as a column vector with n (complex) components, the corresponding bra is a *row* vector having as elements the complex conjugates of the ket elements.

$\langle W|V\rangle = \langle V|W\rangle^*$ then follows automatically from standard matrix multiplication rules, and on multiplying $|V\rangle$ by a complex number a to get $a|V\rangle$ (meaning that each element in the column of numbers is multiplied by a) the corresponding bra goes to $\langle V|a^* = a^*\langle V|$.

But suppose that instead of multiplying a ket by a number, we operate on it with a linear operator.

What generates the parallel transformation among the bras?

In other words, if $A|V\rangle = |V'\rangle$, what operator sends the bra $\langle V|$ to $\langle V'|$?

It must be a linear operator, because A is linear, that is, if under A $|V_1\rangle \rightarrow |V'_1\rangle$, $|V_2\rangle \rightarrow |V'_2\rangle$ and $|V_3\rangle = |V_1\rangle + |V_2\rangle$, then under A $|V_3\rangle$ is required to go to $|V'_3\rangle = |V'_1\rangle + |V'_2\rangle$.

Consequently, under the parallel *bra* transformation we must have $\langle V_1| \rightarrow \langle V'_1|$, $\langle V_2| \rightarrow \langle V'_2|$ and $\langle V_3| \rightarrow \langle V'_3|$, = the bra transformation is necessarily also linear.

Recalling that the bra is an n -element *row* vector, the most general linear transformation sending it to another bra is an $n \times n$ matrix operating on the bra from the *right*.

This bra operator is called the *adjoint* of A , written $A^\dagger$.

That is, the ket $A|V\rangle$ has corresponding bra $\langle V|A^\dagger$.

In an orthonormal basis, using the notation $\langle Ai|$ to denote the bra $\langle i|A^\dagger$ corresponding to the ket $A|i\rangle = |Ai\rangle$, say,

$$(A^\dagger)_{ij} = \langle i|A^\dagger|j\rangle = \langle Ai|j\rangle = \langle j|Ai\rangle^* = A_{ij}^*$$

So the adjoint operator is the transpose complex conjugate.

Important: for a *product* of two operators (prove this!),

$$(AB)^\dagger = B^\dagger A^\dagger$$

An operator equal to its adjoint $A = A^\dagger$ is called *Hermitian*.

As we shall find next, Hermitian operators are of central importance in quantum mechanics.

An operator equal to *minus* its adjoint, $A = -A^\dagger$, is *anti*-Hermitian (sometimes termed skew Hermitian).

These two operator types are essentially generalizations of real and imaginary number: any operator can be expressed as a sum of a Hermitian operator and an anti-Hermitian operator,

$$A = \frac{1}{2}(A + A^\dagger) + \frac{1}{2}(A - A^\dagger)$$

The definition of adjoint naturally extends to *vectors and numbers*: the adjoint of a ket is the corresponding bra, the adjoint of a number is its complex conjugate.

This is useful to bear in mind when taking the adjoint of an operator which may be partially constructed of vectors and numbers, such as projection-type operators.

The adjoint of a product of matrices, vectors and numbers is the product of the adjoints *in reverse order*.

Of course, for numbers the order doesn't matter.

Unitary Operators

An operator is *unitary* if $U^\dagger U = 1$.

This implies first that U operating on any vector gives a vector having the same norm, since the new norm $\langle V | U^\dagger U | V \rangle = \langle V | V \rangle$.

Furthermore, inner products are preserved, $\langle W | U^\dagger U | V \rangle = \langle W | V \rangle$.

Therefore, under a unitary transformation the original orthonormal basis in the space must go to another orthonormal basis.

Conversely, any transformation that takes one orthonormal basis into another one is a unitary transformation.

To see this, suppose that a linear transformation A sends the members of the orthonormal basis $(|1\rangle_1, |2\rangle_1, \dots, |n\rangle_1)$ to the different orthonormal set $(|1\rangle_2, |2\rangle_2, \dots, |n\rangle_2)$, so

$$A |1\rangle_1 = |1\rangle_2, \text{ etc.}$$

Then the vector $|V\rangle = \sum_i v_i |i\rangle_1$ will go to $|V'\rangle = A|V\rangle = \sum_i v_i |i\rangle_2$, having the same norm, $\langle V'|V'\rangle = \langle V|V\rangle = \sum_i |v_i|^2$.

A matrix element $\langle W'|V'\rangle = \langle W|V\rangle = \sum_i w_i^* v_i$, but also $\langle W'|V'\rangle = \langle W|A^\dagger A|V\rangle$.

That is, $\langle W|V\rangle = \langle W|A^\dagger A|V\rangle$ for *arbitrary* kets.

This is only possible if $A^\dagger A = I$, so A is unitary.

A unitary operation amounts to a rotation (possibly combined with a reflection) in the space.

Evidently, since $U^\dagger U = 1$, the adjoint $U^\dagger$ rotates the basis back: it is the inverse operation, and so $UU^\dagger = 1$ also, that is, U and $U^\dagger$ commute.

Determinants

We review in this section the *determinant* of a matrix, a function closely related to the operator properties of the matrix.

Let's start with 2 x 2 matrices:

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

The determinant of this matrix is defined by:


$$\det A = |A| = a_{11}a_{22} - a_{12}a_{21}$$

Writing the two rows of the matrix as vectors:

$$\vec{a}_1^R = (a_{11}, a_{12})$$

$$\vec{a}_2^R = (a_{21}, a_{22})$$

(R denotes row), $\det A = \vec{a}_1^R \times \vec{a}_2^R$ is just the *area* (with appropriate sign) of the parallelogram having the two row vectors as adjacent sides:



This is zero if the two vectors are parallel (linearly dependent) and is not changed by adding any multiple of $\vec{a}_2^R$ to $\vec{a}_1^R$ (because the new parallelogram has the same base and the same height as the original - check this by drawing).

Let's go on to the more interesting case of 3 x 3 matrices:

$$A = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix}$$

The determinant of A is defined as

$$\det A = \varepsilon_{ijk} a_{1i} a_{2j} a_{3k}$$

where ε_{ijk} is 0 if any two are equal, +1 if $ijk = 123, 231$ or 312 (that is to say, an even permutation of 123) and -1 if ijk is an odd permutation of 123.

Repeated suffixes, of course, imply summation here.

Writing this out explicitly,

$$\det A = a_{11}a_{22}a_{33} + a_{21}a_{32}a_{13} + a_{31}a_{12}a_{23} - a_{11}a_{32}a_{23} - a_{21}a_{12}a_{33} - a_{31}a_{22}a_{13}$$


Just as in two dimensions, it's worth looking at this expression in terms of vectors representing the rows of the matrix

$$\vec{a}_1^R = (a_{11}, a_{12}, a_{13})$$

$$\vec{a}_2^R = (a_{21}, a_{22}, a_{23})$$

$$\vec{a}_3^R = (a_{31}, a_{32}, a_{33})$$

so



This is the volume of the parallelopiped formed by the three vectors being adjacent sides (meeting at one corner, the origin).

This parallelepiped volume will of course be zero if the three vectors lie in a plane, and it is not changed if a multiple of one of the vectors is added to another of the vectors.

That is to say, *the determinant of a matrix is not changed if a multiple of one row is added to another row.*

This is because the determinant is linear in the elements of a single row,

$$\det \begin{pmatrix} \vec{a}_1^R + \lambda \vec{a}_2^R \\ \vec{a}_2^R \\ \vec{a}_3^R \end{pmatrix} = \det \begin{pmatrix} \vec{a}_1^R \\ \vec{a}_2^R \\ \vec{a}_3^R \end{pmatrix} + \lambda \det \begin{pmatrix} \vec{a}_2^R \\ \vec{a}_2^R \\ \vec{a}_3^R \end{pmatrix}$$

and the last term is zero because two rows are identical—so the triple vector product vanishes.

A more general way of stating this, applicable to larger determinants, is that for a determinant with two identical rows, the symmetry of the two rows, together with the *antisymmetry* of ε_{ijk} , ensures that the terms in the sum all cancel in pairs.

Since the determinant is not altered by adding some multiple of one row to another, if the rows are linearly dependent, one row could be made identically zero by adding the right multiples of the other rows.

Since every term in the expression for the determinant has one element from each row, the determinant would then be identically zero.

For the three-dimensional case, the linear dependence of the rows means the corresponding vectors lie in a plane, and the parallelepiped is flat.

The algebraic argument generalizes easily to $n \times n$ determinants: they are *identically zero if the rows are linearly dependent*.

The generalization from 3×3 to $n \times n$ determinants is that $\det A = \varepsilon_{ijk} a_{1i} a_{2j} a_{3k}$ becomes:

$$\det A = \varepsilon_{ijk\dots p} a_{1i} a_{2j} a_{3k} \cdots a_{np}$$

where $ijk \dots p$ is summed over all permutations of $123 \dots n$, and the ε symbol is zero if any two of its suffixes are equal, +1 for an even permutation and -1 for an odd permutation.

Note: any permutation can be written as a product of swaps of neighbors.

Such a representation is in general not unique, but for a given permutation, all such representations will have either an odd number of elements or an even number.

An important theorem is that for a product of two matrices A, B the determinant of the product is the product of the determinants, $\det AB = \det A \times \det B$.

This can be verified by brute force for 2×2 matrices, and a proof in the general case can be found in any book on mathematical physics.

It can also be proved that if the *rows* are linearly independent, the determinant cannot be zero.

To return to three dimensions, it is clear from the form of

$$\det A = a_{11}a_{22}a_{33} + a_{21}a_{32}a_{13} + a_{31}a_{12}a_{23} - a_{11}a_{32}a_{23} - a_{21}a_{12}a_{33} - a_{31}a_{22}a_{13}$$

that we could equally have taken the *columns* of A as three vectors, $A = (\vec{a}_1^C, \vec{a}_2^C, \vec{a}_3^C)$ in an obvious notation, $\det A = (\vec{a}_1^C \times \vec{a}_2^C) \cdot \vec{a}_3^C$, and linear dependence among the *columns* will

also ensure the vanishing of the determinant - so, in fact, linear dependence of the columns ensures linear dependence of the rows.

This, too, generalizes to $n \times n$: in the definition of determinant $\det A = \varepsilon_{ijk\dots p} a_{1i} a_{2j} a_{3k} \cdots a_{np}$, the row suffix is fixed and the column suffix goes over all permissible permutations, with the appropriate sign - but the same terms would be generated by having the *column* suffixes kept in numerical order and allowing the row suffix to undergo the permutations.

Eigenkets and Eigenvalues

If an operator A operating on a ket $|V\rangle$ gives a multiple of the same ket,

$$A |V\rangle = \lambda |V\rangle$$

then $|V\rangle$ is said to be an *eigenket* (or, just as often, *eigenvector*, or *eigenstate*!) of A with eigenvalue λ .

Eigenkets and eigenvalues are of central importance in quantum mechanics: dynamical variables are operators, a physical measurement of a dynamical variable yields an eigenvalue of the operator, and forces the system into an eigenket.

In this section, we shall show how to find the eigenvalues and corresponding eigenkets for an operator A .

We'll use the notation $A |a_i\rangle = a_i |a_i\rangle$ for the set of eigenkets $|a_i\rangle$ with corresponding eigenvalues a_i . (Obviously, in the eigenvalue equation here the suffix i is not summed over.)

The first step in solving $A |V\rangle = \lambda |V\rangle$ is to find the allowed eigenvalues a_i .

Writing the equation in matrix form,

$$\begin{pmatrix} A_{11} - \lambda & A_{12} & \cdot & \cdot & A_{1n} \\ A_{21} & A_{22} - \lambda & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ A_{n1} & \cdot & \cdot & \cdot & A_{nn} - \lambda \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ \cdot \\ \cdot \\ v_n \end{pmatrix} = 0$$

This equation is actually telling us that the *columns* of the matrix $A - \lambda I$ are linearly dependent!

To see this, write the matrix as a row vector each element of which is one of its columns, and the equation becomes

$$\left(\vec{M}_1^C, \vec{M}_2^C, \dots, \vec{M}_n^C \right) \begin{pmatrix} v_1 \\ \cdot \\ \cdot \\ \cdot \\ v_n \end{pmatrix} = 0$$

which is to say

$$v_1 \vec{M}_1^C + v_2 \vec{M}_2^C + \dots + v_n \vec{M}_n^C = 0$$

the columns of the matrix are indeed a *linearly dependent* set.

We know that means the determinant of the matrix $A - \lambda I$ is zero,

$$\begin{vmatrix} A_{11} - \lambda & A_{12} & \cdot & \cdot & A_{1n} \\ A_{21} & A_{22} - \lambda & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ A_{n1} & \cdot & \cdot & \cdot & A_{nn} - \lambda \end{vmatrix} = 0$$

Evaluating the determinant using $\det A = \varepsilon_{ijk\dots p} a_{1i} a_{2j} a_{3k} \cdots a_{np}$ gives an n^{th} order polynomial in λ sometimes called the *characteristic polynomial*.

Any polynomial can be written in terms of its roots:

$$C(\lambda - a_1)(\lambda - a_2) \cdots (\lambda - a_n) = 0$$

where the a_i 's are the roots of the polynomial, and C is an overall constant, which from inspection of the determinant we can see to be $(-1)^n$. (It's the coefficient of λ^n .)

The polynomial roots (which we don't yet know) are in fact the eigenvalues.

For example, putting $\lambda = a_1$ in the matrix, $\det(A - a_1 I) = 0$, which means that

$(A - a_1 I) |V\rangle = 0$ has a nontrivial solution $|V\rangle$, and this is our eigenvector $|a_1\rangle$.

Notice that the diagonal term in the determinant $(A_{11} - \lambda)(A_{22} - \lambda) \cdots (A_{nn} - \lambda)$ generates the leading two orders in the polynomial $(-1)^n(\lambda^n - (A_{11} + \cdots + A_{nn})\lambda^{n-1})$, (and some lower order terms too).

Equating the coefficient of λ^{n-1} here with that in $(-1)^{n-1}(\lambda - a_1)(\lambda - a_2) \cdots (\lambda - a_n)$,

$$\sum_{i=1}^n a_i = \sum_{i=1}^n A_{ii} = \text{Tr} A$$

Putting $\lambda = 0$ in both the determinantal and the polynomial representations (in other words, equating the λ -independent terms),

$$\prod_{i=1}^n a_i = \det A$$

So we can find both the sum and the product of the eigenvalues directly from the determinant, and for a 2×2 matrix this is enough to solve the problem.

For anything bigger, the method is to solve the polynomial equation $\det(A - \lambda I) = 0$ to find the set of eigenvalues, then use them to calculate the corresponding eigenvectors.

This is done one at a time.

This looks like n equations for the n numbers v_i , but it isn't: remember the rows are linearly dependent, so there are only $n - 1$ independent equations.

However, that's enough to determine the ratios of the vector components $v_1, \dots, v_n$ then finally the eigenvector is normalized.

The process is then repeated for each eigenvalue.

Extra care is needed if the polynomial has coincident roots.

Eigenvalues and Eigenstates of Hermitian Matrices

For a *Hermitian* matrix, it is easy to establish that the eigenvalues are always *real*.

Note: A basic postulate of Quantum Mechanics, is that physical observables are represented by Hermitian operators.

Taking (in this section) A to be hermitian, $A = A^\dagger$, and labeling the eigenkets by the eigenvalue, that is,

$$A |a_1\rangle = a_1 |a_1\rangle$$

the inner product with the bra $\langle a_1|$ gives $\langle a_1| A |a_1\rangle = a_1 \langle a_1 | a_1\rangle$

But the inner product of the adjoint equation (remembering $A = A^\dagger$)

$$\langle a_1| A = a_1^* \langle a_1|$$

with $|a_1\rangle$ gives $\bar{a}_1 A |a_1\rangle = a_1^* \langle a_1 | a_1\rangle$, so $a_1 = a_1^*$, and all the eigenvalues must be real.

They certainly don't have to all be different: for example, the unit matrix I is Hermitian, and all its eigenvalues are of course 1.

But let's first consider the case where they *are* all different.

It's easy to show that the eigenkets belonging to *different* eigenvalues are *orthogonal*.

If

$$A |a_1\rangle = a_1 |a_1\rangle$$

$$A |a_2\rangle = a_2 |a_2\rangle,$$

take the adjoint of the first equation and then the inner product with $|a_2\rangle$, and compare it with the inner product of the second equation with $\langle a_1|$:

$$\langle a_1|A|a_2\rangle = a_1 \langle a_1|a_2\rangle = a_2 \langle a_1|a_2\rangle$$

so $\langle a_1|a_2\rangle = 0$ unless the eigenvalues are equal.

If they *are* equal, they are referred to as *degenerate* eigenvalues.

Let's first consider the nondegenerate case: A has all eigenvalues distinct.

The eigenkets of A , appropriately normalized, form an orthonormal basis in the space.

Write

$$|a_1\rangle = \begin{pmatrix} v_{11} \\ v_{21} \\ \vdots \\ v_{n1} \end{pmatrix}, \quad \text{and consider the matrix}$$

$$V = \begin{pmatrix} v_{11} & v_{12} & \cdots & v_{1n} \\ v_{21} & v_{22} & \cdots & v_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ v_{n1} & v_{n2} & \cdots & v_{nn} \end{pmatrix} = (|a_1\rangle \quad |a_2\rangle \quad \cdots \quad |a_n\rangle)$$

Now

$$AV = A(|a_1\rangle \quad |a_2\rangle \quad \cdots \quad |a_n\rangle) = (a_1|a_1\rangle \quad a_2|a_2\rangle \quad \cdots \quad a_n|a_n\rangle)$$

so

$$V^\dagger A V = \begin{pmatrix} \langle a_1 | \\ \langle a_2 | \\ \vdots \\ \langle a_n | \end{pmatrix} \begin{pmatrix} a_1 |a_1\rangle & a_2 |a_2\rangle & \cdots & a_n |a_n\rangle \end{pmatrix} = \begin{pmatrix} a_1 & 0 & \cdots & 0 \\ 0 & a_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_n \end{pmatrix}.$$

Note also that, obviously, V is unitary:

$$V^\dagger V = \begin{pmatrix} \langle a_1 | \\ \langle a_2 | \\ \vdots \\ \langle a_n | \end{pmatrix} \begin{pmatrix} |a_1\rangle & |a_2\rangle & \cdots & |a_n\rangle \end{pmatrix} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}$$

We have established, then, that for a Hermitian matrix with distinct eigenvalues (nondegenerate case), the unitary matrix V having columns identical to the normalized eigenkets of A diagonalizes A , that is, $V^\dagger A V$ is diagonal.

Furthermore, its (diagonal) elements equal the corresponding eigenvalues of A .

Another way of saying this is that the unitary matrix V is the transformation from the original orthonormal basis in the space to the basis formed of the normalized eigenkets of A .

Proof that the Eigenvectors of a Hermitian Matrix Span the Space

We'll now move on to the general case: what if some of the eigenvalues of A are the same? In this case, any linear combination of them is also an eigenvector with the same eigenvalue.

Assuming they form a basis in the subspace, the Gram Schmidt procedure can be used to make it orthonormal, and so part of an orthonormal basis of the whole space.

However, we have not actually established that the eigenvectors *do* form a basis in a degenerate subspace.

Could it be that (to take the simplest case) the two eigenvectors for the single eigenvalue turn out to be parallel?

This is actually the case for some 2×2 matrices — for example,

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

we need to prove it is *not* true for Hermitian matrices, and nor are the analogous statements for higher-dimensional degenerate subspaces.

The procedure is by induction from the 2×2 case.

The general 2×2 Hermitian matrix has the form

$$\begin{pmatrix} a & b \\ b^* & c \end{pmatrix}$$

where a, c are real.

It is easy to check that if the eigenvalues are degenerate, this matrix becomes a real multiple of the identity, and so trivially has two orthonormal eigenvectors.

Since we already know that if the eigenvalues of a 2×2 Hermitian matrix are distinct it can be diagonalized by the unitary transformation formed from its orthonormal eigenvectors, we have established that *any* 2×2 Hermitian matrix can be so diagonalized.

To carry out the induction process, we now assume any Hermitian matrix can be diagonalized by a unitary transformation.

We need to prove this means it's also true for an $n \times n$ Hermitian matrix A .

Recall a unitary transformation takes one complete orthonormal basis to another.

If it diagonalizes a Hermitian matrix, the new basis is necessarily the set of orthonormalized eigenvectors.

Hence, if the matrix can be diagonalized, the eigenvectors do span the n -dimensional space.

Choose an eigenvalue a_1 of A , with normalized eigenvector $|a_1\rangle = (v_{11}, v_{21}, \dots, v_{n1})^T$.

We put in T for transpose, to save the awkwardness of filling the page with a few column vectors.

We construct a unitary operator V by making this the first column, then filling in with $n-1$ other normalized vectors to construct, with $|a_1\rangle$, an n -dimensional orthonormal basis.

Now, since $A|a_1\rangle = a_1|a_1\rangle$, the first column of the matrix AV will just be $a_1|a_1\rangle$, and the rows of the matrix $V^\dagger = V^{-1}$ will be $\langle a_1|$ followed by $n-1$ normalized vectors orthogonal to it, so the first column of the matrix $V^\dagger AV$ will be a_1 followed by zeros.

It is easy to check that $V^\dagger A V$ is Hermitian, since A is, so its first row is also zero beyond the first diagonal term.

This establishes that for an $n \times n$ Hermitian matrix, a unitary transformation exists to put it in the form:

$$V^\dagger A V = \begin{pmatrix} a_1 & 0 & \cdot & \cdot & 0 \\ 0 & M_{22} & \cdot & \cdot & M_{2n} \\ 0 & \cdot & \cdot & \cdot & \cdot \\ 0 & \cdot & \cdot & \cdot & \cdot \\ 0 & M_{n2} & \cdot & \cdot & M_{nn} \end{pmatrix}.$$

But we can now perform a second unitary transformation in the $(n - 1) \times (n - 1)$ subspace orthogonal to $|a_1\rangle$ (this of course leaves $|a_1\rangle$ invariant), to complete the full diagonalization -- that is to say, the existence of the $(n - 1) \times (n - 1)$ diagonalization, plus the argument above, guarantees the existence of the $n \times n$ diagonalization: the induction is complete.

Diagonalizing a Hermitian Matrix

As discussed above, a Hermitian matrix is diagonal in the orthonormal basis of its set of eigenvectors: $|a_1\rangle, |a_2\rangle, \dots, |a_n\rangle$, since

$$\langle a_i | A | a_j \rangle = \langle a_i | a_j | a_j \rangle = a_j \langle a_i | a_j \rangle = a_j \delta_{ij}.$$

If we are given the matrix elements of A in some other orthonormal basis, to diagonalize it we need to rotate from the initial orthonormal basis to one made up of the eigenkets of A .

Denoting the initial orthonormal basis in the standard fashion

$$|1\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad |2\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad |i\rangle = \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix} \dots (1 \text{ in } i^{\text{th}} \text{ place down}), \quad |n\rangle = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}.$$

the elements of the matrix are $A_{ij} = \langle i | A | j \rangle$.

$$|V\rangle \rightarrow |V'\rangle = U |V\rangle.$$

Under this transformation, the matrix element

$$\langle W | A | V \rangle \rightarrow \langle W' | A | V' \rangle = \langle W | U^\dagger A U | V \rangle.$$

So we can find the appropriate transformation matrix U by requiring that $U^\dagger A U$ be diagonal with respect to the *original* set of basis vectors.

Transforming the operator in this way, leaving the vector space alone, is equivalent to rotating the vector space and leaving the operator alone.

Of course, in a system with more than one operator, the same transformation would have to be applied to all the operators.

In fact, just as we discussed for the nondegenerate (distinct eigenvalues) case, the unitary matrix U we need is just composed of the normalized eigenkets of the operator A ,

$$U = (|a_1\rangle, |a_2\rangle, \dots, |a_n\rangle)$$

and it follows as before that

$$(U^\dagger A U)_{ij} = \langle a_i | a_j | a_j \rangle = \delta_{ij} a_j, \text{ a diagonal matrix.}$$

The repeated suffixes here are of course *not* summed over.

If some of the eigenvalues are the same, the Gram Schmidt procedure may be needed to generate an orthogonal set, as mentioned earlier.

Functions of Matrices

The same unitary operator U that diagonalizes an Hermitian matrix A will also diagonalize A^2 because

$$U^{-1} A^2 U = U^{-1} A A U = U^{-1} A U U^{-1} A U$$

so

$$U^\dagger A^2 U = \begin{pmatrix} a_1^2 & 0 & 0 & \cdot & 0 \\ 0 & a_2^2 & 0 & \cdot & 0 \\ 0 & 0 & a_3^2 & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & \cdot & \cdot & \cdot & a_n^2 \end{pmatrix}.$$

Evidently, this same process works for any power of A , and formally for any function of A expressible as a power series, but of course convergence properties need to be considered, and this becomes trickier on going from finite matrices to operators on infinite spaces.

Commuting Hermitian Matrices

From the above, the set of powers of an Hermitian matrix all commute with each other, and have a common set of eigenvectors (but not the same eigenvalues, obviously).

In fact it is not difficult to show that any two Hermitian matrices that commute with each other have the same set of eigenvectors (after possible Gram Schmidt rearrangements in degenerate subspaces).

If two $n \times n$ Hermitian matrices A, B commute, that is, $AB = BA$, and A has a *nondegenerate* set of eigenvectors $A|a_i\rangle = a_i|a_i\rangle$, then $AB|a_i\rangle = BA|a_i\rangle = Ba_i|a_i\rangle = a_iB|a_i\rangle$, that is, $B|a_i\rangle$ is an eigenvector of A with eigenvalue a_i .

Since A is nondegenerate, $B|a_i\rangle$ must be some multiple of $|a_i\rangle$, and we conclude that A, B have the same set of eigenvectors.

Now suppose A is degenerate, and consider the $m \times m$ subspace S_{a_i} spanned by the eigenvectors $|a_i, 1\rangle, |a_i, 2\rangle, \dots$ of A having eigenvalue a_i .

Applying the argument in the paragraph above, $B|a_i, 1\rangle, B|a_i, 2\rangle, \dots$ must also lie in this subspace.

Therefore, if we transform B with the same unitary transformation that diagonalized A , B will not in general be diagonal in the subspace S_{a_i} , but it will be what is termed *block diagonal*, in that if B operates on any vector in S_{a_i} it gives a vector in S_{a_i} .

B can be written as two diagonal blocks: one $m \times m$, one $(n - m) \times (n - m)$, with zeroes outside these diagonal blocks, for example, for $m = 2, n = 5$:

$$\begin{pmatrix} b_{11} & b_{12} & 0 & 0 & 0 \\ b_{21} & b_{22} & 0 & 0 & 0 \\ 0 & 0 & b_{33} & b_{34} & b_{35} \\ 0 & 0 & b_{43} & b_{44} & b_{45} \\ 0 & 0 & b_{53} & b_{54} & b_{55} \end{pmatrix}$$

And, in fact, if there is only *one* degenerate eigenvalue that second block will only have nonzero terms on the diagonal:

$$\begin{pmatrix} b_{11} & b_{12} & 0 & 0 & 0 \\ b_{21} & b_{22} & 0 & 0 & 0 \\ 0 & 0 & b_3 & 0 & 0 \\ 0 & 0 & 0 & b_4 & 0 \\ 0 & 0 & 0 & 0 & b_5 \end{pmatrix}$$

B therefore operates on two subspaces, one m -dimensional, one $(n - m)$ -dimensional, *independently* — a vector entirely in one subspace stays there.

This means we can complete the diagonalization of B with a unitary operator that *only* operates on the $m \times m$ block S_{a_i} .

Such an operator will also affect the eigenvectors of A , but that doesn't matter, because all vectors in this subspace are eigenvectors of A with the *same* eigenvalue, so as far as A is concerned, we can choose any orthonormal basis we like — the basis vectors will still be eigenvectors.

This establishes that any two *commuting* Hermitian matrices can be diagonalized at the same time.

Obviously, this can never be true of *noncommuting* matrices, since all diagonal matrices commute.

Diagonalizing a Unitary Matrix

Any unitary matrix can be diagonalized by a unitary transformation.

To see this, recall that *any* matrix M can be written as a sum of a Hermitian matrix and an anti-Hermitian matrix,

$$M = \frac{M+M^\dagger}{2} + \frac{M-M^\dagger}{2} = A + iB$$

where both A , B are Hermitian.

This is the matrix analogue of writing an arbitrary complex number as a sum of real and imaginary parts.

If A, B commute, they can be simultaneously diagonalized (see the previous section), and therefore M can be diagonalized. Now, if a *unitary* matrix is expressed in this form $U = A + iB$ with A, B Hermitian, it easily follows from $UU^\dagger = U^\dagger U = 1$ that A, B commute, so *any unitary matrix U can be diagonalized by a unitary transformation.*

More generally, if a matrix M commutes with its adjoint $M^\dagger$, it can be diagonalized.

Note: it is *not* possible to diagonalize M unless both A and B are simultaneously diagonalized.

This follows from $U^\dagger AU, U^\dagger iBU$ being Hermitian and anti-Hermitian for any unitary operator U , so their off-diagonal elements cannot cancel each other, they must all be zero if M has been diagonalized by U , in which case the two transformed matrices $U^\dagger AU, U^\dagger iBU$ are diagonal, therefore commute, and so do the original matrices A, B .

It is worthwhile looking at a specific example, a simple rotation of one orthonormal basis into another in three dimensions.

Obviously, the axis through the origin about which the basis is rotated is an eigenvector of the transformation.

It's less clear what the other two eigenvectors might be — or, equivalently, what are the eigenvectors corresponding to a two dimensional rotation of basis in a plane?

The way to find out is to write down the matrix and diagonalize it.

The matrix

$$U(\theta) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}.$$

Note that the determinant is equal to unity.

The eigenvalues are given by solving

$$\begin{vmatrix} \cos \theta - \lambda & \sin \theta \\ -\sin \theta & \cos \theta - \lambda \end{vmatrix} = 0 \text{ to give } \lambda = e^{\pm i\theta}.$$

The corresponding eigenvectors satisfy

$$\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} u_1^\pm \\ u_2^\pm \end{pmatrix} = e^{\pm i\theta} \begin{pmatrix} u_1^\pm \\ u_2^\pm \end{pmatrix}.$$

The eigenvectors, normalized, are:

$$\begin{pmatrix} u_1^\pm \\ u_2^\pm \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ \pm i \end{pmatrix}.$$

Note that, in contrast to a Hermitian matrix, the eigenvalues of a unitary matrix do not have to be real.

In fact, from $U^\dagger U = 1$, sandwiched between the bra and ket of an eigenvector, we see that any eigenvalue of a unitary matrix must have unit modulus — it's a complex number on the unit circle.

With hindsight, we should have realized that one eigenvalue of a two-dimensional rotation had to be $e^{i\theta}$, the product of two two-dimensional rotations is given by adding the angles of rotation, and a rotation through π changes all signs, so has eigenvalue -1.

Note that the eigenvector itself is independent of the angle of rotation: in two dimensions, the rotations all commute, so they must have common eigenvectors.

Successive rotation operators applied to the plus eigenvector add their angles, when applied to the minus eigenvector, all angles are subtracted.