

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation, or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

ARAM HARROW: So let's get started. This week Professor Zwiebach is away, and I'll be doing today's lecture. And Will Detmold will do the one on Wednesday.

The normal office hours, unfortunately, will not be held today. One of us will cover his hours on Wednesday though. And you should also just email either me or Professor Detmold if you want to set up an appointment to talk to in the next few days.

What I'm going to talk about today will be more about the linear algebra that's behind all of quantum mechanics. And, at the end of last time-- last lecture you heard about vector spaces from a more abstract perspective than the usual vectors are columns of numbers perspective. Today we're going to look at operators, which act on vector spaces, which are linear maps from a vector space to itself.

And they're, in a sense, equivalent to the familiar idea of matrices, which are squares or rectangles of numbers. But they work in this more abstract setting of vector spaces, which has a number of advantages.

For example, of being able to deal with infinite dimensional vector spaces and also of being able to talk about basis independent properties. And so I'll tell you all about that today. So we'll talk about how to define operators, some examples, some of their properties, and then finally how to relate them to the familiar idea of matrices.

I'll then talk about eigenvectors and eigenvalues from this operator perspective. And, depending on time today, a little bit about inner products, which you'll hear more about in the future. These numbers here correspond to the sections of the notes that these refer to.

So let me first-- this is a little bit mathematical and perhaps dry at first. The payoff is more distant than usual for things you'll hear in quantum mechanics. I just want to mention a little bit about the motivation for it.

So operators, of course, are how we define observables. And so if we want to know what the properties of observables, of which a key example are of Hamiltonians, then we need to know about operators.

They also, as you will see in the future, are useful for talking about states. Right now, states are described as elements of a vector space, but in the future you'll learn a different formalism in which states are also described as dense operators. What are called density operators or density matrices.

And finally, operators are also useful in describing symmetries of quantum systems. So already in classical mechanics, symmetries have been very important for understanding things like momentum conservation and energy conservation so on. They'll be even more important in quantum mechanics and will be understood through the formalism of operators.

So these are not things that I will talk about today but are sort of the motivation for understanding very well the structure of operators now.

So at the end of the last lecture, Professor Zwiebach defined linear maps. So this is the set of linear maps from a vector space, v , to a vector space w .

And just to remind you what it means for a map to be linear, so T is linear if for all pairs of vectors in v , the way T acts on their sum is given by just T of u plus T of v . That's the first property.

And second, for all vectors u and for all scalars a -- so f is the field that we're working over, it could be reals or complexes-- we have that if T acts on a times u , that's equal to a times T acting on u .

So if you put these together what this means is that T essentially looks like multiplication. The way T acts on vectors is precisely what you would expect from

the multiplication map, right? It has the distributive property and it commutes with scalars.

So this is sort of informal-- I mean, the formal definition is here, but the informal idea is that T acts like multiplication. So if the map that squares every entry of a vector does not act like this, but linear operators do.

And for this reason we often neglect the parentheses. So we just write Tu to mean T of u , which is justified because of this analogy with multiplication.

So an important special case of this is when v is equal to w . And so we just write I of v to denote the maps from v to itself. Which you could also write like this. And these are called operators on v . So when we talk about operators on a vector space, v , we mean linear maps from that vector space to itself.

So let me illustrate this with a few examples. Starting with some of the examples of vector spaces that you saw from last time. So one example of a vector space is an example you've seen before but a different notation. This is the vector space of all real polynomials in one variable.

So real polynomials over some variable, x . And over-- this is an infinite dimensional vector space-- and we can define various operators over it. For example, we can define one operator, T , to be like differentiation.

So what you might write as $\frac{d}{dx}$, and it's defined for any polynomial, p , to map p to p' . So this is certainly a function from polynomials to polynomials.

And you can check that it's also linear if you multiply the polynomial by a scalar, then the derivative multiplied by the same scale. If I take the derivative of a sum of two polynomials, then I get the sum of the derivatives of those polynomials. I won't write that down, but you can check that the properties are true. And this is indeed a linear operator.

Another operator, which you've seen before, is multiplication by x . So this is defined as the map that simply multiplies the polynomial by x . Of course, this gives you

another polynomial. And, again, you can check easily that it satisfies these two conditions.

So this gives you a sense of why things that don't appear to be matrix-like can still be viewed in this operator picture. Another example, which you'll see later shows some of the slightly paradoxical features of infinite dimensional vector spaces, comes from the vector space of infinite sequences.

So these are all the infinite sequences of reals or complexes or whatever f is. One operator we can define is the left shift operator, which is simply defined by shifting this entire infinite sequence left by one place and throwing away the first position.

So you start with x_2, x_3 , and so. Still goes to infinity so it still gives you an infinite sequence. So it is indeed a map-- that's the first thing you should check that this is indeed a map from V to itself-- and you can also check that it's linear, that it satisfies these two products.

Another example is right shift. And here-- Yeah?

AUDIENCE: So left shift was the first one or--

ARAM HARROW: That's right. So there's no back, really. It's a good point. So you'd like to not throw out the first one, perhaps, but there's no canonical place to put it in. This just goes off to infinity and just falls off the edge.

It's a little bit like differentiation. Right?

AUDIENCE: Yeah. I guess it loses some information.

ARAM HARROW: It loses some information. That's right.

It's a little bit weird, right? Because how many numbers do you have before you applied the left shift? Infinity. How many do you have after you applied the left shift? Infinity. But you lost some information. So you have to be a little careful with the infinities. OK

The right shift. Here it's not so obvious what to do. We've kind of made space for another number, and so we have to put something in that first position. So this will be question mark x_1 , x_2 , dot, dot, dot.

Any guesses what should go in the question mark?

AUDIENCE: 0?

ARAM HARROW: 0. Right. And why should that be 0?

AUDIENCE: [INAUDIBLE].

ARAM HARROW: What's that?

AUDIENCE: So it's linear.

ARAM HARROW: Otherwise it wouldn't be linear. Right.

So imagine what happens if you apply the right shift to the all 0 string. If you were to get something non-zero here, then you would map to the 0 vector to a non-zero vector.

But, by linearity, that's impossible. Because I could take any vector and multiply it by the scalar 0 and I get the vector 0. And that should be equal to the scalar 0 multiplied by the output of it. And so that means that T should always map 0 to 0. T should always map the vector 0 to the vector 0. And so if we want right shift to be a linear operator, we have to put a 0 in there.

And this one is strange also because it creates more space but still preserves all of the information.

So two other small examples of linear operators that come up very often. There's, of course, the 0 operator, which takes any vector to the 0 vector. Here I'm not distinguishing between-- here the 0 means an operator, here it means a vector. I guess I can clarify it that way. And this is, of course, linear and sends any vector space to itself.

One important thing is that the output doesn't have to be the entire vector space. The fact that it sends a vector space to itself only means that the output is contained within the vector space. It could be something as boring as 0 that just sends all the vectors to a single point.

And finally, one other important operator is the identity operator that sends-- actually I won't use the arrows here. We'll get used to the mathematical way of writing it-- that sends any vector to itself.

Those are a few examples of operators. I guess you've seen already kind of the more familiar matrix-type of operators, but these show you also the range of what is possible.

So the space L of v of all operators-- I want to talk now about its properties. So L of v is the space of all linear maps from v to itself. So this is the space of maps on a vector space, but itself is also a vector space.

So the set of operators satisfies all the axioms of a vector space. It contains a 0 operator. That's this one right here.

It's closed under a linear combination. If I add together two linear operators, I get another linear operator.

It's closed under a scalar multiplication. If I multiply a linear operator by a scalar, I get another linear operator, et cetera.

And so everything we can do on a vector space, like finding a basis and so on, we can do for the space of linear operators. However, in addition to having the vector space structure, it has an additional structure, which is multiplication.

And here we're finally making use of the fact that we're talking about linear maps from a vector space to itself. If we were talking about maps from v to w , we couldn't necessarily multiply them by other maps from v to w , we could only multiply them by maps from w to something else. Just like how, if you're multiplying rectangular matrices, the multiplication is not always defined if the dimensions don't match up,

But since these operators are like square matrices, multiplication is always defined, and this can be used to prove many nice things about them. So this type of structure being a vector space of multiplication makes it, in many ways, like a field-- like real numbers or complexes-- but without all of the properties.

So the properties that the multiplication does have first is that it's associative. So let's see what this looks like. So if we have a times bc is equal to ab times c .

And the way we can check this is just by verifying the action of this on any vector. So an operator is defined by its action and all of the vectors in a vector space. So the definition of ab can be thought of as asking how does it act on all the possible vectors? And this is defined just in terms of the action of a and b as you first apply b and then you apply A .

So this can be thought of as the definition of how to multiply operators. And then from this, you can easily check the associativity property that in both cases, however you write it out, you obtain A of B of C of v .

I'm writing out all the parentheses just to sort of emphasize this is C acting on v , and then B acting on C of v , and then A acting on all of this. The fact that this is equal-- that this is the same no matter how A , B , and C are grouped is again part of what let's us justify this right here, where we drop-- we just don't use parentheses when we have operators acting.

So, yes, we have the associative property. Another property of multiplication that operators satisfy is the existence of an identity. That's just the identity operator, here, which for any vector space can always be defined.

But there are other properties of multiplication that it doesn't have. So inverses are not always defined. They sometimes are. I can't say that a matrix is never invertible, but for things like the reals and the complexes, every nonzero element has an inverse. And for matrices, that's not true.

And another property-- a more interesting one that these lack-- is that the

multiplication is not commutative. So this is something that you've seen for matrices. If you multiply two matrices, the order matters, and so it's not surprising that same is true for operators.

And just to give a quick example of that, let's look at this example one here with polynomials. And let's consider S times T acting on the monomial x to the n .

So T is differentiation so it sends this to n times x to the n minus 1. So we get S times n , x to the n minus 1. Linearity means we can move the n past the S . S acting here multiplies by x , and so we get n times x to the n .

Whereas if we did the other order, we get T times S acting on x to the n , which is x to the n plus 1. When you differentiate this you get n plus 1 times x to the n .

So these numbers are different meaning that S and T do not commute. And it's kind of cute to measure to what extent do they not commute. This is done by the commutator. And what these equations say is that if the commutator acts on x to the n , then you get n plus 1 times x to the n minus n times x to the n , which is just x to the n .

And we can write this another way as identity times x to the n . And since this is true for any choice of n , it's true for what turns out to be a basis for the space of polynomials. So $1x$, x squared, x cubed, et cetera, these span the space of polynomials.

So if you know what an operator does and all of the x to the n 's, you know what it does on all the polynomials. And so this means, actually, that the commutator of these two is the identity.

And so the significance of this is-- well, I won't dwell on the physical significance of this, but it's related to what you've seen for position and momentum. And essentially the fact that these don't commute is actually an important feature of the theory.

So these are some of the key properties of the space of operators. I want to also now tell you about some of the key properties of individual operators. And basically,

if you're given an operator and want to know the gross features of it, what should you look at?

So one of these things is the null space of an operator. So this is the set of all v , of all vectors, that are killed by the operator. They're sent to 0.

In some case-- so this will always include the vector 0. So this always at least includes the vector 0, but in some cases it will be a lot bigger. So for the identity operator, the null space is only the vector 0. The only thing that gets sent to 0 is 0 itself. Whereas, for the 0 operator, everything gets sent to 0. So the null space is the entire vector space.

For left shift, the null space is only 0 itself-- sorry, for right shift the null space is only 0 itself. And what about for left shift? What's the null space here? Yeah?

AUDIENCE: Some number with a string of 0s following it.

ARAM HARROW: Right. Any sequence where the first number is arbitrary, but everything after the first number is 0. And so from all of these examples you might guess that this is a linear subspace, because in every case it's been a vector space, and, in fact, this is correct.

So this is a subspace of v because, if there's a vector that gets sent to 0, any multiple of it also will be sent to 0. And of the two vectors that get sent to 0, their sum will also be sent to 0. So the fact that it's a linear subspace can be a helpful way of understanding this set.

And it's related to the properties of T as a function. So for a function we often want to know whether it's 1 to 1, or injective, or whether it's [? onto ?] or surjective.

And you can check that if T is injective, meaning that if u is not equal to v , then T of u is not equal to T of v . So this property, that T maps distinct vectors to distinct vectors, turns out to be equivalent to the null space being only the 0 vector.

So why is that? This statement here, that whenever u is not equal to v , T of u is not equal to T of v , another way to write that is whenever u is not equal to v , T of u

minus v is not equal to 0 .

And if you look at this statement a little more carefully, you'll realize that all we cared about on both sides was u minus v . Here, obviously, we care about u minus v . Here we only care if u is not equal to v .

So that's the same as saying if u minus v is non-zero, then T of u minus v is non-zero. And this in turn is equivalent to saying that the null space of T is only 0 .

In other words, the set of vectors that get sent to 0 consists only of the 0 vector itself. So the null space for linear operators is how we can characterize whether they're 1 to 1, whether they destroy any information.

The other subspace that will be important that we will use is the range of an operator. So the range of an operator, which we can also just write as T of v , is the set of all points that vectors in v get mapped to. So the set of all Tv for some vector, v .

So this, too, can be shown to be a subspace. And that's because-- it takes a little more work to show it, but not very much-- if there's something in the output of T , then whatever the corresponding input is we could have multiplied that by a scalar. And then the corresponding output also would get multiplied by a scalar, and so that, too, would be in the range.

And so that means that for anything in the range, we can multiply it by any scalar and again get something in the range. Similarly for addition. A similar argument shows that the range is closed under addition. So indeed, it's a linear subspace. Again, since it's a linear subspace, it always contains 0 . And depending on the operator, may contain a lot more.

So whereas the null space determined whether T was injective, the range determines whether T is surjective. So the range of T equals v if and only if T is surjective.

And here this is simply the definition of being surjective. It's not really a theorem like

it was in the case of T being injective. Here that's really what it means to be surjective is that your output is the entire space.

So one important property of the range of the null space whenever v is finite dimensional is that the dimension of v is equal to the dimension of the null space plus the dimension of the range. And this is actually not trivial to prove. And I'm actually not going to prove it right now. But the intuition of it is as follows.

Imagine that v is some n dimensional space and the null space has dimension k . So that means you have input of n degrees of freedom, but T kills k of n . And so k at different degrees of freedom, no matter how you vary them, have no effect on the output. They just get mapped to 0.

And so what's left are n minus k degrees of freedom that do affect the outcome. Where, if you vary them, it does change the output in some way. And those correspond to the n minus k dimensions of the range.

And if you want to get formal, you have to formalize what I was saying about what's left is n minus k . If you talk about something like the orthogonal complement or completing a basis or in some way formalize that intuition.

And, in fact, you can do a little further, and you can decompose the space. So this is just dimensions counting. You can even decompose the space into the null space and the complement of that and show that T is 1 to 1 on the complement of the null space.

But for now, I think this is all that we'll need for now. Any questions so far? Yeah?

AUDIENCE: Why isn't the null space part of the range?

ARAM HARROW: Why isn't it part of the range?

AUDIENCE: So you're taking T of v and null space is just the special case when T of v is equal to 0.

ARAM HARROW: Right. So the null space are all of the-- This theorem, I guess, would be a little bit

more surprising if you realized that it works not only for operators, but for general linear maps. And in that case, the range is a subset space of w . Because the range is about the output. And the null space is a [? subset ?] space of v , which is part of the input.

And so in that case, they're not even comparable. The vectors might just have different lengths. And so it can never-- like the null space in a range, in that case, would live in totally different spaces.

So let me give you a very simple example. Let's suppose that T is equal to $\begin{pmatrix} 3 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$. So just a diagonal 4 by 4 matrix. Then the null space would be the span of e_2 , that's the vector with a 1 in the second position. And the range would be the span of e_1, e_3, e_4 .

So in fact, usually it's the opposite that happens. The null space in the range are-- in this case they're actually orthogonal subspaces. But this picture is actually a little bit deceptive in how nice it is.

So if you look at this, total space is 4, four dimensions, it divides up into one dimension that gets killed, and three dimensions where the output still tells you something about the input, where there's some variation of the output.

But this picture makes it seem-- the simplicity of this picture does not always exist. A much more horrible example is this matrix. So what's the null space of this matrix? Yeah?

AUDIENCE: You just don't care about the upper [INAUDIBLE].

ARAM HARROW: You don't care about the-- informally, it's everything of this form. Everything with something in the first position, 0 in the second position. In other words, it's the span of e_1 .

What about the range?

AUDIENCE: [INAUDIBLE].

ARAM HARROW: What's that? Yeah?

AUDIENCE: [INAUDIBLE].

ARAM HARROW: It's actually--

AUDIENCE: Isn't it e1?

ARAM HARROW: It's also e1. It's the same thing.

So you have this intuition that some degrees of freedom are preserved and some are killed. And here they look totally different. And there they look the same. So you should be a little bit nervous about trying to apply that intuition.

You should be reassured that at least the theorem is still true. At least 1 plus 1 is equal to 2. We still have that.

But the null space and the range are the same thing here. And the way around that paradox-- Yeah?

AUDIENCE: So can you just change the basis-- is there always a way of changing the basis of the matrix? In this case it becomes [INAUDIBLE]? Or not necessarily?

ARAM HARROW: No. It turns out that, even with the change of basis, you cannot guarantee that the null space and the range will be perpendicular. Yeah?

AUDIENCE: What if you reduce it to only measures on the-- or what if you reduce the matrix of-- [? usability ?] on only [INAUDIBLE] on the diagonal?

ARAM HARROW: Right Good. So if you do that, then-- if you do row [? education ?] with two different row and column operations, then what you've done is you have a different input and output basis. And so that would-- then once you kind of unpack what's going on in terms of the basis, then it would turn out that you could still have strange behavior like this.

What your intuition is based on is that if the matrix is diagonal in some basis, then you don't have this trouble. But the problem is that not all matrices can be

diagonalized. Yeah?

AUDIENCE: So is it just the trouble that the null is what you're acting on and the range is what results from it?

ARAM HARROW: Exactly. And they could even live in different space. And so they really just don't-- to compare them is dangerous.

So it turns out that the degrees of freedom corresponding to the range-- what you should think about are the degrees of freedom that get sent to the range. And in this case, that would be e_2 .

And so then you can say that e_1 gets sent to 0 and e_2 gets sent to the range. And now you really have decomposed the input space into two orthogonal parts. And because we're talking about a single space, the input space, it actually makes sense to break it up into these parts.

Whereas here, they look like they're the same, but really input and output spaces you should think of as potentially different. So this is just a mild warning about reading too much into this formula, even though it's the rough idea it counting degrees of freedom is still roughly accurate.

So I want to say one more thing about properties of operators, which is about invertibility. And maybe I'll leave this up for now.

So we say that a linear operator, T , has a left inverse, S , if multiplying T on the left by S will give you the identity. And T has a right inverse, S' , you can guess what will happen here if multiplying T on the right by S' gives you identity.

And what if T has both? Then in that next case, it turns out that S and S' have to be the same. So here's the proof.

So if both exist, then S is equal to S' times identity-- by the definition of the identity. And then we can replace identity with TS' . Then we can group these and cancel them and get S' .

So if a matrix has both a left and a right inverse, then it turns out that the left and right inverse are the same. And in this case, we say that T is invertible, and we define T inverse to be S .

One question that you often want to ask is when do left to right inverses exist? Actually, maybe I'll write it here.

Intuitively, there should exist a left inverse when, after we've applied T , we haven't done irreparable damage. So whatever we're left with, there's still enough information that some linear operator can restore our original vector and give us back the identity.

And so that condition is when-- of not doing irreparable damage, of not losing information, is asking essentially whether T is injective. So there exists a left inverse if and only if T is injective.

Now for a right inverse the situation is sort of dual to this. And here what we want-- we can multiply on the right by whatever we like, but there won't be anything on the left. So after the action of T , there won't be any further room to explore the whole vector space.

So the output of T had better cover all of the possibilities if we want to be able to achieve identity by multiplying T by something on the right.

So any guesses for what the condition is for having a right inverse?

AUDIENCE: Surjective.

ARAM HARROW: Surjective. Right. So there exists a right inverse if and only if T is surjective.

Technically, I've only proved one direction. My hand waving just now proved that, if T is not injective, there's no way it will have a left inverse. If it's not surjective, there's no way it'll have a right inverse.

I haven't actually proved that, if it is injective, there is such a left inverse. And if it is surjective, there is such a right universe. But those I think are good exercises for

you to do to make sure you understand what's going on.

This takes us part of the way there. In some cases our lives become much easier. In particular, if v is finite dimensional, it turns out that all of these are equivalent. So T is injective if and only if T is surjective if and only if T is invertible.

And why is this? Why should it be true that T is surjective if and only if T is injective? Why should those be equivalent statements? Yeah?

AUDIENCE: This isn't really a rigorous statement, but if the intuition of it is a little bit that you're taking vectors in v to vectors in v .

ARAM HARROW: Yeah.

AUDIENCE: And so your mapping is 1 to 1 if and only if every vector is mapped to, because then you're not leaving anything out.

ARAM HARROW: That's right. In failing to be injective and failing to be surjective both look like losing information. Failing to be injective means I'm sending a whole non-zero vector and its multiples to 0, that's a degree of freedom lost. Failing to be surjective means once I look at all the degrees of freedom I reach, I haven't reached everything. So they intuitively look the same.

So that's the right intuition. There's a proof, actually, that makes use of something on a current blackboard though. Yeah?

AUDIENCE: Well, you need the dimensions of-- so if the [INAUDIBLE] space is 0, you need dimensions of [? the range to p. ?]

ARAM HARROW: Right. Right. So from this dimensions formula you immediately get because if this is 0, then this is the whole vector space. And if this is non-zero, this is not the whole vector space.

And this proof is sort of non-illuminating if you don't know the proof of that thing-- which I apologize for. But also, you can see immediately from that that we've used the fact that v is finite dimensional.

And it turns out this equivalence breaks down if the vector space is infinite dimensional. Which is pretty weird.

There's a lot of subtleties of infinite dimensional vector spaces that it's easy to overlook if you build up your intuition from matrices.

So does anyone have an idea of a-- so let's think of an example of a vector of something that is on an infinite dimensional space that's surjective but not injective. Any guesses for such an operation? Yeah?

AUDIENCE: The left shift.

ARAM HARROW: Yes. You'll notice I didn't erase this blackboard strategically. Yes. The left shift operator is surjective. I can prepare any vector here I like just by putting it into the $x_2, x_3, \text{dot, dot, dot}$ parts.

So the range is everything, but it's not injective because it throws away the first register. It maps things with a non-zero element in the first position and 0's everywhere else to 0. So this is surjective not injective.

On the other hand, if you want something that's injective and not surjective, you don't have to look very far, the right shift is injective and not surjective.

It's pretty obvious it's not surjective. There's that 0 there which definitely means it cannot achieve any vector. And it's not too hard to see it's injective. It hasn't lost any information. It's like you're in the hotel that's infinitely long and all the rooms are full and the person at the front desk says, no problem. I'll just move everyone down one room to the right, and you can take the first room.

So that policy is injective-- you'll always get a room to yourself-- and made possible by having an infinite dimensional vector space.

So in infinite dimensions we cannot say this. Instead, we can say that T is invertible if and only if T is injective and surjective. So this statement is true in general for infinite dimensional, whatever, vector spaces. And only in the nice special case of

finite dimensions do we get this equivalence. Yeah?

AUDIENCE: Can the range and null space of T a [INAUDIBLE] of T the operator again use a vector space [INAUDIBLE]?

ARAM HARROW: Yes. the question was do the null space in a range are they properties just of T or also of v ? And definitely you also need to know v . The way I've been writing it, T is implicitly defined in terms of v , which in turn is implicitly defined in terms of the field, f . And all these things can make a difference. Yes?

AUDIENCE: So do you have to be a bijection for it to be--

ARAM HARROW: That's right. That's right. Invertible is the same a bijection.

So let me now try and relate this to matrices. I've been saying that operators are like the fancy mathematician's form of matrices. If you're Arrested Development fans, it's like magic trick versus an illusion. But are they different or not depends on your perspective. There are advantages to seeing it both ways, I think.

So let me tell you how you can view an operator in a matrix form. The way to do this-- and the reason why matrices are not universally loved by mathematicians is I haven't specified a basis this whole time. But if I want a matrix, all I needed was a vector space and a function-- a linear function between two vector spaces-- or, sorry, from a vector space to itself.

But if I want a matrix, I need additional structure. And mathematicians try to avoid that whenever possible. But if you're willing to take this additional structure-- so if you choose a basis v_1 through v_n -- it turns out you can get a simpler form of the operator that's useful to compute with.

So why is that? Well, the fact that it's a basis that means that any v can be written as linear combinations of these basis elements where a_1 though an belong to the field. And since T is linear, if T acts on v , we can rewrite it in this way, and you see that the entire action is determined by T acting on v_1 through v_n .

So think about-- if you wanted to represent an operator in a computer, you'd say,

well, there's an infinite number of input vectors. And for each input vector I have to write down the output vector. And this says, no, you don't. You only need to restore on your computer what does T do to v_1 , what does T do to v_2 , et cetera.

So that's good. Now you only have to write down n vectors, and since these factors in turn can be expressed in terms of the basis, you can express this just in terms of a bunch of numbers.

So let's further expand Tv_j in this basis. And so there's some coefficient. So it's something times v_1 plus something times v_2 something times v_n . And I'm going to-- these somethings are a function of T so I'm just going to call this T_{1j} , T_{2j} , T_{nj} . And this whole thing I can write more succinctly in this way.

And now all I need are these T's of ij , and that can completely determine for me the action of T because this Tv here-- so Tv we can write as a sum over j of T times $a_j v_j$. And we can move the a_j past the T. And then if we expand this out, we get that it's a sum over i from 1 to n , sum over j from 1 to n , of $T_{ij} a_j v_i$.

And so if we act on in general vector, v , and we know the coefficients of v in some basis, then we can re-express it in that basis as follows. And this output in general can always be written in the basis with some coefficients. So we could always write it like this.

And this formula tells you what those coefficients should be. They say, if your input vector has coefficients a_1 through a_n , then your output vector has coefficients b_1 through b_n , where the $b_{sub i}$ are defined by this sum.

And of course there's a more-- this formula is one that you've seen before, and it's often written in this more familiar form. So this is now the familiar matrix-vector multiplication. And it says that the b vector is obtained from the a vector by multiplying it by the matrix of these T_{ij} .

And so this T is the matrix form-- this is a matrix form of the operator T. And you might find this not very impressive. You say, well, look I already knew how to

multiply a matrix by vector. But what I think is nice about this is that the usual way you learn linear algebra if someone says, a vector is a list of numbers. A matrix is a rectangle of numbers. Here's are the rules for what you do with them. If you want to put them together, you do it in this way.

Here this was not an axiom of the theory at all. We just started with linear maps from one vector space to another one and the idea of a basis as something that you can prove has to exist and you can derive matrix multiplication. So matrix multiplication emerges-- or matrix-vector multiplication emerges as a consequence of the theory rather than as something that you have to put in.

So that, I think, is what's kind of cute about this even if it comes back on the end to something that you had been taught before. Any questions about that?

So this is matrix-vector multiplication. You can similarly derive matrix-matrix multiplication.

So if we have two operators, T and S , and we act on a vector, v_k -- and by what I argued before, it's enough just to know how they act on the basis vectors. You don't need to know-- and once you do that, you can figure out how they act on any vector.

So if we just expand out what we wrote before, this is equal to T times the sum over j of $S_{jk} v_j$. So Sv_k can be re-expressed in terms of the basis with some coefficients. And those coefficients will depend on the vector you start with, k , and the part of the basis that you're using to express it with j .

Then we apply the same thing again with T . We get-- this is sum over i , sum over j $T_{ij} S_{jk} v_i$.

And now, what have we done? TS is an operator and when you act of v_k it spat out something that's a linear combination of all the basis states, v_i , and the coefficient of v_i is this part in the parentheses.

And so this is the matrix element of TS . So the ik matrix element of ts is the sum

over j of Tijssk. And so just like we derived matrix-vector multiplication, here we can derive matrix-matrix multiplication.

And so what was originally just sort of an axiom of the theory is now the only possible way it could be if you want to define operator multiplication is first one operator acts, than the other operator acts.

So in terms of this-- so this, I think, justifies why you can think of matrices as a faithful representation of operators. And once you've chosen a basis, they can-- the square full of numbers becomes equivalent to the abstract map between vector spaces.

And the equivalent-- they're so equivalent that I'm just going to write things like equal signs. Like I'll write identity equals a bunch of 1's down the diagonal, right? And not worry about the fact that technically this is an operator and this is a matrix.

And similarly, the 0 matrix equals a matrix full of 0's. Technically, we should write-- if you want to express the basis dependence, you can write things like T parentheses-- sorry, let me write it like this.

If you really want to be very explicit about the basis, you could use this to refer to the matrix. Just to really emphasize that the matrix depends not only on the operator, but also on your choice of basis. But we'll almost never bothered to do this. We usually just sort of say it in words what the basis is.

So matrices are an important calculational tool, and we ultimately want to compute numbers of physical quantities so we cannot always spend our lives in abstract vector spaces.

But the basis dependence is an unfortunate thing. A basis is like a choice of coordinate systems, and you really don't want your physics to depend on it, and you don't want quantity if you compute to be dependent on.

And so we often want to formulate-- we're interested in quantities that are basis independent. And in fact, that's a big point of the whole operator picture is that

because the quantities we want are ultimately basis independent, it's nice to have language that is itself basis independent. Terminology and theorems that do not refer to a basis.

I'll mention a few basis independent quantities, and I won't say too much more about them because you will prove properties [INAUDIBLE] on your p set, but one of them is the trace and another one is the determinant.

And when you first look at them-- OK, you can check that each one is basis independent, and it really looks kind of mysterious. I mean, like, who pulled these out of the hat? They look totally different, right? They don't look remotely related to each other. And are these all there is? Are there many more?

And it turns out that, at least for matrices with eigenvalues, these can be seen as members of a much larger family. And the reason is that the trace turns out to be the sum of all the eigenvalues and the determinant turns out to be the product of all of the eigenvalues. And in general, we'll see in a minute, that basis independent things-- actually, not in a minute. In a future lecture-- that basis independent things are functions of eigenvalues.

And furthermore, that don't care about the ordering of the eigenvalues. So they're symmetric functions of eigenvalues. And then it starts to make a little bit more sense. Because if you talk about symmetric polynomials, those are two of the most important ones where you just add up all the things and when you multiply all the things. And then, if you add this perspective of symmetric polynomial of the eigenvalue, then you can cook up other basis independent quantities.

So this is actually not the approach you should take on the p set. The [? p set ?] asks you to prove more directly that the trace is basis independent, but the sort of framework that these fit into is symmetric functions of eigenvalues.

So I want to say a little bit about eigenvalues. Any questions about matrices before I do?

So eigenvalues-- I guess, these are basis independent quantities. Another important

basis independent quantity, or property of a matrix, is its eigenvalue-eigenvector structure.

The place where eigenvectors come from is by considering a slightly more general thing, which is the idea of an invariant subspace. So we say that U is a T invariant subspace if T of U -- this is an operator acting on an entire subspace. So what do I mean by that? I mean the set of all TU for vectors in the subspace. If T of U is contained in U .

So I take a vector in this subspace, act on it with T , and then I'm still in the subspace no matter which vector I had. So some examples that always work. The 0 subspace is invariant. T always maps it to itself. And the entire space, v , T is a linear operator on v so by definition it maps v to itself.

These are called the trivial examples. And usually when people talk about non-trivial invariant subspaces they mean not one of these two. The particular type that we will be interested in are one dimensional ones.

So this corresponds to a direction that T fixes. So U -- this vector space now can be written just as the span of a single vector, U , and U being T invariant is equivalent to TU being a μU , because they're just a single vector. So all I have to do is get that single vector right and I'll get the whole subspace right.

And that, in turn, is equivalent to TU being some multiple of U . And this equation you've seen before. This is the familiar eigenvector equation. And if it's a very, very important equation it might be named after a mathematician, but this one is so important that two of the pieces of it have their own special name. So these are called-- λ is called an eigenvalue and U is called an eigenvector.

And more or less it's true that all of the solutions to this are called eigenvalues, and all the solutions are called eigenvectors. There's one exception, which is there's one kind of trivial solution to this equation, which is when U is 0 this equation is always true. And that's not very interesting, but it's true for all values of λ .

And so that doesn't count as being an eigenvalue. And you can tell a doesn't

correspond to 1D invariant subspace, right? It corresponds to a 0 dimensional subspace, which is the trivial case.

So we say that λ is an eigenvalue of T if Tu equals λU for some non-zero vector, U . So the non 0 is crucial. And then the spectrum of T is the collection of all eigenvalues.

So there's something a little bit asymmetric about this, which is we still say that 0 vector is an eigenvector with all the various eigenvalues, but we had to put this here or everything would be an eigenvalue and it wouldn't be very interesting. So the--

Oh, also I want to say this term spectrum you'll see it other [INAUDIBLE]. You'll see spectral theory or spectral this or that, and that means essentially making use of the eigenvalues. So people talk about partitioning a graph using eigenvalues of the associated matrix, that's called spectral partitioning. And so throughout math, this term is used a lot.

So I have only about three minutes left to tell-- so I think I will not finish the eigenvalue discussion but will just show you a few examples of how it's not always as nice as you might expect.

So one example that I'll consider is the vector space will be the reals, 3D real space, and the operator, T , will be rotation about the z-axis by some small angle. Let's call it a θ rotation about the z-axis.

Turns out, if you write this in matrix form, it looks like this. $\cos \theta$ $\sin \theta$ 0 $-\sin \theta$ $\cos \theta$ 0 0 0 1 . That 1 is because it leaves the z-axis alone and then x and y get rotated.

You can tell if θ is 0 it does nothing so that's reassuring. And if θ does a little bit, then it starts mixing the x and y components. So that is the rotation matrix.

So what is an eigenvalue-- and anyone say what an eigenvalue is of this matrix?

AUDIENCE: 1.

ARAM HARROW: 1. Good. And what's the eigenvector?

AUDIENCE: The z basis vector.

ARAM HARROW: The z basis vector. Right. So it fixes a z basis vector so this is an eigenvector with eigenvalue 1.

Does it have any other eigenvectors? Yeah?

AUDIENCE: If you go to the complex plane, then yes.

ARAM HARROW: If you are talking about complex numbers, then yes, it has complex eigenvalues. But if we're talking about a real vector space, then it doesn't. And so this just has one eigenvalue and one eigenvector.

And if we were to get rid of the third dimension-- so if we just had T-- and let's be even simpler, let's just take theta to be pi over 2. So let's just take a 90 degree rotation in the plane.

Now T has no eigenvalues. There are no vectors other than 0 that it sends to itself. And so this is a slightly unfortunate note to end the lecture on. You think, well, these eigenvalues are great, but maybe they exist, maybe they don't.

And you'll see next time part of the reason why we use complex numbers, even though it looks like real space isn't complex, is because any polynomial can be completely factored in complex numbers, and every matrix has a complex eigenvalue.

OK, I'll stop here.