HERBERT GROSS: Hi. Last time, you may recall that we were looking at some of our examples in three-dimensional space and talking about picking a couple of vectors in three-dimensional space, be they i and j or alpha and beta. And we then talked about the subspaces spanned by alpha and beta or the subspace span by i and j, noticing that the subspace could be smaller than the entire space. Or possibly, it could even equal the entire space. In no event could it be greater than the entire space.

At any rate, what we would like to do today is to generalize this idea into more than two or three dimensions. And also, to make the generalization go independently of our n-tuple notation that we used earlier in the course. At any rate, today's lesson is entitled, Spanning Vectors, and the idea is this.

Let's suppose we have a vector space, v, and we now arbitrarily choose n vectors, alpha 1 up to alpha n and v. For example, in the lecture last time, some of our examples involved v being three-dimensional space and n being two. In other words, we picked i and j, alpha and beta, in three-dimensional space.

But at any rate, what we do is this. We take alpha 1 up to alpha n, any n arbitrarily chosen, vectors, n, v, and we look at the set, which I write s of alpha 1 up to alpha n. We look at the set of all linear combinations of alpha 1 up to alpha n.

And written compactly, that's written this way. In other words, it's a c1 alpha 1 plus, et cetera, cn alpha n. Where the cs are real numbers, they belong to r.

Now, my claim is that this set is more than a set. It's a space. Namely, if you of two linear combinations of alpha 1 up to alpha n, you get a linear combination of alpha 1 up to alpha n. If you take a scalar multiple of a linear combination of alpha 1 up to alpha n, you, again, get a linear combination of alpha 1 up to alpha n.

Algebraically, this is quite simple, because structurally we mimic what's happening in ordinary algebra. And what we're saying is, you can add the two alpha 1 components together, the alpha 2 components together, et cetera. I'll illustrate that in a simple example in a moment, but the idea is this.

What we say is, by definition-- and we've done this in some of the exercises already, you'll recall-- that the subspace consisting of all linear combinations of alpha 1 up to alpha n, which is a subspace of v, it's called the space spanned by alpha 1 up to alpha n. That's all this thing means. The space spanned by alpha 1 up to alpha n means the set of all linear combinations of alpha 1 up to alpha n, and that is a subspace of v.

For example, the easiest case I can think of is to pick n equals 1, given a vector space, v, simply pick any elements, say alpha 1, that belongs to v. And all linear combinations of alpha 1 are simply what? All scalar multiples of alpha 1. In other words, the space spanned by alpha 1 is just a set of all constant multiples of alpha. So c alpha 1, where c is a real number.

Now, my claim is that, clearly, the sum of any two scalar multiples of alpha 1 is a scalar multiple of alpha 1. And a scalar multiple of a scalar multiple of alpha 1 is, again, a scalar multiple of alpha 1. I'd like to show you that, though, independently of a geometric interpretation. This follows clearly from our axiomatic approach, namely, suppose beta and gamma are any two elements in the space spanned by alpha 1.

That means beta is a scalar multiple as alpha 1, say c1 alpha 1. Gamma is a scalar multiple of alpha 1, say c2 alpha 1. Therefore, beta plus gamma is c1 alpha 1 plus c2 alpha 1.

By our distributive property, that's c1 plus c2 times alpha 1. But since c1 and c2 a real numbers, and the sum of two real numbers is, again, a real number, we see that beta plus gamma is a real number times alpha 1, in other words, a scalar multiple of alpha 1. And that means that by definition it belongs to the space spanned by alpha 1.

Similarly, if we take beta and multiply it by the real number r, r times c1 alpha 1, by the associated property for scalar multiplication, is the scale of multiple r times c1 times alpha 1. And if r and c1 are real number, so is their product. So consequently, r times beta is a scalar times alpha 1, which by definition means it's in the space spanned by alpha 1.

And geometrically, if you want to see what this thing means, notice that the space spanned by alpha 1-- if you want think of alpha 1 as an arrow-- the space spanned by alpha 1 is simply what? The set of all linear combinations of alpha-- all scalar multiples of alpha 1, and that is precisely the straight line determined by alpha 1. In other words, any linear combination-- any scalar multiple of alpha 1 has what?

The same direction as alpha 1. In other words, lies in the same direction, but may have a different magnitude and/or a different sense. But certainly, the line part doesn't change.

Now, the key point is this, and this is going to play a very key role throughout the remainder of our discussion. One would believe that the more vectors you take linear combinations of, the bigger your space becomes. In other words, if you take the space spanned by two vectors, then tack onto those two vectors a third vector, you might believe that the resulting space must be larger than the original space.

Well, obviously, it can't be any smaller, and I'll explain what obviously means in the moment, if it's not that obvious. But the key point is, that if I take not just alpha 1 up to alpha n, but an additional vector, say, alpha 1 up to alpha n plus the additional vector alpha n plus 1, and look at the space spanned by alpha 1 through alpha n plus 1, that space need not be larger than the space spanned by alpha 1 up to alpha n. In other words, merely by tacking on a new vector, you do not necessarily tack on to the space spanned by the vectors.

By means of a trivial example, let alpha 1 be i, let alpha 2 be j, an alpha 3 be i plus j. Clearly, the space spanned by alpha 1 and alpha 2 is by definition this sort of all linear combinations of alpha 1 alpha 2. And in terms of this simple example, that obviously is the x, y plane, because alpha 1 is i, alpha 2 as j. Now, let's look at the space spanned by alpha 1, alpha 2, and alpha 3.

Again, geometrically, notice that that space is still the x, y plane. Because after all, alpha 1, alpha 2, and alpha 3, being i, j, and i plus j, respectively, all lie in the x, y plane. Consequently, any linear combination of them will lie in the x, y plane. But let's pretend we didn't know that.

Let's, first of all, observe that from the pure mathematics of this alone, that the space spanned by alpha 1 alpha 2 must be a subspace of the space spanned by alpha 1, alpha 2, an alpha 3. And the proof, by the way, is quite trivial. Namely, by definition, the space spanned by alpha 1, alpha 2, alpha 3 is a set of all linear combinations of alpha 1, alpha 2, alpha 3. In particular, if I choose c3 to equal 0, notice that this term drops out. And in particular, what I'm left with is c1 alpha 1 plus c2 alpha 2. In other words, every linear combination of alpha 1 and alpha 2 is, in particular, a linear combination of alpha 1, alpha 2, and alpha 3.

But my claim is, just as we saw geometrically-- let me now prove this algebraically-- that in particular, if beta is any vector belonging to the space spanned by alpha 1, alpha 2, and alpha 3, axiomatically, it must belong to the space spanned by alpha 1 and alpha 2. We see that geometrically. We already know that the space is still the x, y plane. How would we prove this axiomatically?

Well, if beta belongs to this space, by definition it's a linear combination as alpha 1, alpha 2, alpha 3. Noticing, that alpha 3 is by definition alpha 1 plus alpha 2. I can replace alpha 3 by alpha 1 plus alpha 2.

I can now distribute c3 with alpha 1 and alpha 2. I can now combine the alpha 1 terms and the alpha 2 terms together. In other words, c1 plus c3 alpha 1 plus c2 plus c3 alpha 2. Notice, that since c1, c2, and c3 are real numbers, c1 plus c3 is the real number. c2 plus c3 is a real number.

Consequently, beta is a linear combination of alpha 1 and alpha 2. Therefore, beta is in the space spanned by alpha 1 and alpha 2. Therefore, the space spanned by alpha 1, alpha 2, and alpha 3 is a subspace of the space spanned by alpha 1 and alpha 2. And since each of the subspaces is a subspace of the other, we see that these two spaces happen to be equal in this sense.

And by the way, if you're looking for a very quick way of seeing how this transcends such a simple application, notice that when we were dealing with differential equations, we looked for solutions that were linearly independent. In other words, we wanted solutions that could not be expressed as linear combinations of the previous solutions. And the reason for that was, if they could be expressed this way, the arbitrary constants weren't independent in the sense that they could be amalgamated. At any rate, whichever interpretation you want here, notice that this leads to a generalized definition. Namely, given the n vectors, alpha 1 up to alpha n and v, they are called linearly dependent.

Now, this may look like a messy expression that might frighten you. All this says is they're called linearly dependent if one of the vectors can be written as a linear combination of the others. And by the way, without proving it here, but leaving that for the exercises, it can be shown that not only when this can be done, but when it can be done, the vector can be written as a linear combination of the vectors that came before. In other words, if any vector, when you list these in any order, if any one of the vectors in the given order can be expressed in terms of a linear combination of the preceding ones, then the set is called linearly dependent. If this can't happen, then they are called linearly independent.

Now, the reason I give this definition first-- it's a bad definition, because it seems to depend on the order in which I list the vectors. The reason I give this definition first is because, I think, intuitively, it's easier to see. In a moment, I will give you a different definition that's equivalent to this one. But first, let me exploit this fairly intuitive definition.

By means of an example, suppose I'm given the four vectors, i plus j, i minus j, 5i plus j, and i plus j plus k. My claim is that these four vectors form a linearly-dependent set, because, in particular, in the given order, 5i plus j can be expressed as a linear combination of i plus j and i minus j. Namely, 5i plus j is three times the first vector here, i plus j, plus 2 times the second vector, i minus j.

By the way, in the next lecture, we'll show how you pick off these coefficients very quickly, but we'll let that go until next time. For now, all that's important is notice that this vector is a linear combination of these two. Notice, by the way, you could have said, but couldn't you express i minus j as a linear combination of i plus j and 5i plus j?

The answer is, yes. In other words, if I had written the 5i plus j and the i minus j in an interchanged order, notice that it would be i minus j that would have been a linear combination of i plus j and 5i plus j. That's why this is kind of a hazy definition. It seems to depend on the order in which the vectors are written.

A more conventional definition is the following. And just to show you that dependence and independence are themselves independent, meaning you can start with either one, let me give you an equivalent definition, but now stressing linear independence instead. Namely, notice that, if I pick all 0 coefficients, in other words, 0 alpha 1 plus a set with 0 alpha n, obviously, that linear combination will add up to 0. If that's the only way you can make a linear combination add up to 0, then the vectors are said to be linearly independent.

In other words, alpha 1 up to alpha n are called linearly independent if the only way that a linear combination of them can equal 0 is if every coefficient is 0. Again, keep in mind that the converse is trivial. If every coefficient is 0, then, obviously, this will be 0. For linearly-independent vectors what we're saying is, the only way this can be 0 is if all the coefficients are 0.

And again, to emphasize linear dependence, what that would mean is what? That the vectors would be linearly dependent if you could find a linear combination, c1 alpha 1 plus, et cetera, cn alpha n equal to 0, where at least one of the cs was not equal to 0. Now, I'm going to stress that in the exercises. This stuff is so new to many of you, I believe, that it's very, very difficult to teach it in great depth during a lecture. I think it has to be kept on a fairly superficial level, leaving additional details to the exercises.

Rather than hammer home these exercises now, let's simply point out that linear dependence implies redundancy. In other words, if one of the vectors can be expressed as linear combination of the others, it can be deleted when you're talking about a set of spanning vectors. In other words, going back to our example of the space spanned by alpha 1 up to alpha n, there are no redundancies if the alphas happen to be independent. But if they're dependent, it means that any vector which can be expressed as a linear combination of the others can be deleted without losing anything in the space.

And I'll show you that in terms of other examples later. But let me show what that means from another point of view also that I want to bring out later in the lecture, and that is another key point. That if the alphas happen to be linearly independent, then if beta belongs to the space spanned by the alphas, not only can beta be written as a linear combination of the alphas, but it can be done so in one and only one way. In other words, to say this in a different form, if the alphas are linearly independent, if beta belongs to the space spanned by the alphas, it has a unique representation as a linear combination of the alphas.

And what that means from another point of view is, that if the alphas are linearly independent, the only way that two linear combinations can be equal is if they are equal coefficient by coefficient. By the way, that was a technique that we used in undetermined coefficients. And again, that shows you why functions, like the exponential, sine, and cosine, et cetera, tie-in with this generalized definition of a vector space. That all of these things have applications far beyond any interpretation due to arrows alone.

But the proof of something like this, given that these two linear combinations are equal, simply transpose and collect like terms, which we can do algebraically by the axioms that we're permitted to work with here. If we group these terms, we get this. Since the alphas are linearly independent, the only way a linear combination can be 0 is if each of the coefficients is 0. That means, in particular, that a1 minus b1 must be 0, et cetera, up to an minus bn being 0. And that says, in addition, what?

That a1 must equal b1, et cetera, an equals bn, and that verifies the claim, that if the alphas are independent, linearly independent, then two different sets of coefficients must lead to two different vectors. In terms of example two, you see, notice that alpha 3 was a linear combination of alpha 1 and alpha 2. Remember, alpha 3 was alpha 1 plus alpha 2. Notice that, therefore, I can express alpha 3 in more than one way as a linear combination of alpha 1, alpha 2, and alpha 3.

For example, certainly alpha 3 is 0 alpha 1 plus 0 alpha 2 plus 1 alpha 3. That's just alpha 3. On the other hand, since alpha 1 plus alpha 2 is alpha 3, and 1 alpha 1 is alpha 1 and 1 alpha 2 is alpha 2, notice that alpha 3 can also be written as 1 alpha 1 plus 1 alpha 2 plus 0 alpha 3. See, the alphas were linearly dependent in this example.

Notice that this allows me to express alpha 3 as two different linear combinations of alpha 1, alpha 2, and alpha 3. By the way, it's not just these two. There are instantly many different ways I can do this. Namely, since alpha 3 is alpha 1 plus alpha 2, pick c to be an arbitrary constant, and look at the expression c alpha 1 plus c alpha 2 times 1 minus c alpha 3.

If I expand this, this becomes c alpha 1 plus c alpha 2 minus c alpha 3. That's minus c alpha 1 plus alpha 2. That cancels, and all I have left is alpha 3. In other words, alpha 3 can be written in infinitely many ways as a linear combination of alpha 1, alpha 2, and alpha 3, because the alphas were linearly dependent. At any rate, again, let me emphasize, I'll drill on this with the exercises.

What I'd like to do now is revisit the concept of dimension, which we treated in terms of n-tuples in terms of spanning vectors and linear independence. Namely, let's suppose I have any vector space now, not necessarily one written in terms of n-tuples. Let's suppose that this other vectors and the 0 vector in the vector space, because, after all, if the only vector in the vector space is the 0 vector, it's not a very exciting vector space.

So let me assume that I have something other than a 0 vector in here. Let me pick any non-zero vector in v. Let me call it alpha 1. And I now-- look at the space-- span by alpha 1. Clearly, that's a subspace of v.

On the other hand-- or not on the other hand. Well, on the other hand, this might be all of v. If the space spanned by alpha 1 is all of v, we stop at that point and say, v is one-dimensional, and it's spanned by the single vector, alpha 1. On the other hand, suppose s, the space spanned by alpha 1, it's not all of v. Well, what does that mean?

It means there must be a vector, say alpha 2, which belongs to v, but not to the space spanned by alpha 1. Well, what we do now is augment alpha 1 by alpha 2. And we now look at the space spanned by alpha 1 and alpha 2.

The first thing I claim is that the space spanned by alpha 1 and alpha 2 is a proper superspace of the space spanned by alpha 1. In other words, it's a space spanned by alpha 1. It's contained in but not equal to the space spanned by alpha 1 and alpha 2. Why is this?

Well, for one thing, alpha 2 by definition didn't belong to the space spanned by alpha 1. See, we picked alpha 2 not to be in the space spanned by alpha 1. On the other hand, alpha 2 does belong to the space spanned by alpha 1 and alpha 2.

You see, notice that the space spanned by alpha 1 and alpha 2 consists of all linear combinations of alpha 1 and alpha 2, and in particular, alpha 2 can be written as 0 alpha 1 plus 1 alpha 2, which makes it a linear combination of alpha 1 and alpha 2. So alpha 2 does not belong to the space spanned by alpha 1, but it does belong to the space spanned by alpha 1 and alpha 2.

Notice, therefore, that we have now enlarged our space. Notice also that alpha 1 alpha 2 must be linearly independent. In other words, we chose any alpha 2, provided only that it belonged to v but not to the space spanned by alpha 1.

But that alpha 2, no matter how we choose it, must be linearly independent of alpha 1, because, after all, if they were linearly dependent, that would mean that alpha 2 is a scalar multiple as alpha 1. But that's a contradiction, because if alpha 2 were a scalar multiple of alpha 1, by definition it would belong to the space spanned by alpha 1. But we specifically chose it not to belong to that space.

At any rate, we continue on in this way. Namely, the space spanned by alpha 1 and alpha 2 might be all of v, in which case we then say v is two-dimensional and that it's spanned by alpha 1 and alpha 2. Or it may be that the space spanned by alpha 1 and alpha 2 is a proper subset of v and doesn't give me all of v.

In this event, I now pick an alpha 3, which belongs to v, but not to the space spanned by alpha 1 and alpha 2. Again, notice that alpha 1, alpha 2, and alpha 3 must be linearly independent since if alpha 3 could be expressed as a linear-- see, we know alpha 2 can't be expressed as a linear combination of alpha 1, because alpha 2 was linearly independent of alpha 1. Now we tack on alpha 3.

The only possibility is is alpha 3 a linear combination of alpha 1 and alpha 2? And the answer is, it couldn't be. Because if alpha 3 were a linear combination of alpha 1 and alpha 2, by definition it would belong to the space spanned by alpha 1 and alpha 2.

But by construction, alpha 3 was chosen not to be in that space. Consequently, if alpha 3 were in that space, that would obviously be a contradiction. In other words, alpha 3 cannot be a linear combination of alpha 1 and alpha 2. Therefore, these three vectors are linearly independent. And similarly, alpha 3, not belonging to the space spanned by alpha 1 and alpha 2, enlarges the space spanned by alpha 1 and alpha 2.

In terms of a rough picture here, what we're saying is, we start with some vector, alpha 1. s of alpha 1 is some subspace. We put an alpha 2, which doesn't belong to the space spanned by alpha 1.

Then the space spanned by alpha 1 and alpha 2 is this larger subspace, which contains the space spanned by alpha 1. If that isn't all of v, we pick a third vector, alpha 3, which isn't in the space spanned by alpha 1 and alpha 2. And we then form the space spanned by alpha 1, alpha 2, and alpha 3, which will contain the space spanned by alpha 1 and alpha 2 as a subspace.

And to make a long story short, we simply continue in this way until we find alpha 1, say, up to alpha n. In other words, we continue for n steps this way, such that v will eventually be the space spanned by alpha 1 up to alpha n. And let me emphasize that this need not happen. In other words, v might be such a huge vector space, that continuing in this way, we could go on forever and not exhaust all of v this way.

But suppose, for the sake of argument, that this does happen, that we find n alphas such that v is spanned by alpha 1 up to alpha n, where the alphas are constructed, as I've indicated previously. In this case, we say that v has dimension n written this way, and the set of vectors, alpha 1 up to alpha n, is called a basis for v. And the reason that it's called the basis for v is that not only can every vector be written as a linear combination of the alphas, every vector in v written as a linear combination of the alphas, but it has a unique representation in terms of the alphas, because the alphas were linearly independent.

In other words, once you have a basis, you see, any vector in v can be written uniquely as an n-tuple where the coefficients stand for the coefficients of alpha 1 up to alpha n. You see, in other words, each v and v has a unique representation, unique representation, as a linear combination of alpha 1 up to alpha n. By the way, if this process doesn't come to an end, then v is set to be an infinite-dimensional vector space.

Now, what I'd like to do is take a few minutes to summarize this particular lecture. Because what we're going to be doing between now and the next lecture is going to involve a lot of work on your part. And that is, that in going through this overview of what you mean by dimension, and spanning sets, and what have you, and linear dependence and independence, many subtleties have occurred. For example, if I have an n-dimensional vector space going through this bit of constructing the space spanned by alpha-- alpha 1 and alpha 2, et cetera-- questions come up.

What if I had started with different vectors? What if I had started with beta 1 instead of alpha 1, and then picked a new vector, beta 2 instead of alpha 2? Could I have spanned v in fewer than n vectors?

In other words, does the dimension depend on the process by which I span the space by how I choose the vectors? Obviously, if this happens to be true, we're in dire trouble. Because we would like to believe intuitively, based on our old definition of dimension, that the dimension of a vector space should not depend on the representative vectors that you pick. This is what's bad about the n-tuple representation.

So you see, what we're going to do in the exercises is show what other ingredients come up in dimension and what we have to worry about. In our next lecture, we will summarize the results that we will prove in the exercises. So you can begin the next lecture even if you're a little bit hazy on some of the exercises. And then we will show, from a practical point of view, once the theory is developed, how do we actually go about constructing a basis.

How do we find what a dimension is in real life? How do we figure out how to construct the alphas if we're not given the abstract situation mentioned in this particular example? At any rate, all I want to do now is review the lecture, do the exercises, start to feel at home with the concepts, and we will drill more on this and review in our next lecture. At any rate, then, until next time. Goodbye.

Funding for the publication of this video was provided by the Gabriella and Paul Rosenbaum Foundation. Help OCW continue to provide free and open access to MIT courses by making a donation at ocw.mit.edu/donate.