MITOCW | Part III: Linear Algebra, Lec 3: Constructing Bases

HERBERT GROSS: Hi. In our last two lectures, we've been talking quite theoretically about what we mean by spanning vectors, dimension, vector spaces. Today, what we'd like to do is add some computational know-how to our bag of knowledge. And let me say at the outset by way of a reinforcement of what we've already said in the study guide, much of the material in this block is very difficult if you're not used to it.

And we have already reminded you, but let's do it again, watch the lectures, get as much out of them as you can, do the exercises, which hopefully will cement down many aspects of the lecture. And then if certain things still aren't crystal clear, watch the lecture a second time. Hopefully, you may not even have to watch the lecture the second time, but try not to read too much into the lectures.

This is material, which I think that once you get the first insight, you must then really work hard and do a lot of exercises. But at any rate, let's go into today's lesson, which I call, Constructing Bases. And by way of review, let's see what theory we've had going into today's lesson.

First of all, given a vector space, v, with certain elements, alpha 1 up to alpha n, the set of n elements, alpha 1 up to alpha n, is called a basis for v, provided two things happen. First, the alphas must span v. And secondly, the alphas must be linearly independent.

And, again, if this gives you any trouble, think in terms of the naive example of the plane or three-dimensional space. For example, in three-dimensional space, i, j, and k span three-dimensional space. And they also happen to be linearly independent.

At any rate, what we then proved was, that if alpha 1 up to alpha n is a basis for v-- and don't be blinded by the alphas, now. What we're really saying here is, that if you have found a basis for v, which has n elements, then the following things must be true. One, any set of more than n elements, for example, any n plus 1 elements of v, are linearly dependent.

Again, in terms of three-dimensional space, what we're saying is, if you have four or more vectors in three space, those vectors may or may not span v. But the thing that we're sure of is that they must be linearly dependent. You can't have more than three linearly-independent vectors in three-dimensional space.

Secondly, fewer than n elements cannot span v. Again, in terms of our three-dimensional example, no two vectors in three space, no two vectors in three space, can span three space. In other words, then, a basis has two properties.

If you have too many vectors, you cannot have linear independence, and if you have too few, you can't span the space. So what is it, then, that a basis must have? If one basis has n elements, every basis must have n elements.

Again, by way of review, why? More than n couldn't be linearly independent. Fewer than n could not span the space. So every basis for v has n elements once we know that one basis has n elements.

Now, this, of course, does not mean that any set of n elements will be a basis, again, in terms of three-dimensional space. Suppose, for example, that I take the following three vectors of three-dimensional space. Oh, heck, let's say i, j, and i plus j. Those are three vectors, i, j, and i plus j, but they only span two-dimensional space. The trouble with i, j, and i plus j is that they were linearly dependent.

You see, all we know is that every basis for v has n elements. It does not mean that every set of n elements will be a basis for v. However, what is true is, that if we have a set of n linearly-independent elements, or else a set of n elements which span v, then if either of those two criteria are obeyed, then the set of n elements will be a basis for v. Again, in three-dimensional space, if you have three linearly-independent vectors in three space, they will be a basis.

They will automatically span the space. And if you have three vectors which span the space in three space, they must be linearly independent. That's all we're saying.

In other words, to say that the dimensions of v equals n is unambiguous. In other words, if we define the dimension of a space to be the number of elements in any basis, then what we're saying is that that is unambiguous. That once one basis has n elements, all bases will have n elements. And now, with respect to a particular basis, what we're saying this.

Suppose we know the dimensions of v is n, and we have a particular basis for v, say, u1 up to un, a particular basis that we single out, then with respect to that basis, each vector, v, in the vector space, v, has a unique representation as a linear combination of the us. In other words, if the us are a basis for v, a particular vector in v can be written in one and only one way, as a linear combination of the us. Which means, you see, that relative to the us, v may be written unambiguously as the n-tuple c1 up to cn, where that's an abbreviation for c1, u1, plus, et cetera, c1, un.

And in this case, we also write that v is equal to-- and to indicate that we're talking about the n-dimensional space, v, with respect to a particular basis, u1 up to un, we simply enclose the set of basis vectors in square brackets. In other words, this expression is read, v is being viewed as the n-dimensional space that has u1 up to un as its basis. And unless otherwise specified, every time we write an n-tuple in v, it's understood that that n-tuple is being abbreviated relative to the basis u1 up to un.

And I think, again, the best way to illustrate this is by means of an example. And, in fact, the remainder of today's lesson will be one example wherein we will take one problem and try to look at it from as many different points of view as possible so that we get a broad spectrum as to what's involved in constructing bases, looking at n-tuples, and how the n-tuple depends on the basis, et cetera. So the example I have in mind is the following.

Let's suppose that. All I know is that I have a four-dimensional vector space. What does that mean again?

It means if the dimensions is four, that any basis has four elements. Let me now pick a particular basis, which I will call u1, u2, u3 and u4. And then notice that notation now. I'm saying that v is the vector space, which has u1, u2, u3, and u4 as a basis. What that means is, if I now elect to use the 4-tuple notation for v, that unless anything is said to the contrary, it means that that 4-tuple is relative to the vectors u1, u2, u3, and u4.

For example, let's suppose I was to say to you, describe the subspace of v spanned by the four vectors alpha 1, alpha 2, alpha 3, alpha 4, where alpha 1 is the 4-tuple 1, 1, 2, 3, alpha 2 is the 4-tuple 2, 3, 4, 5, alpha 3 is the 4-tuple 3, 7, 6, 5, and alpha 4 is the 4-tuple 4, 5, 9, 9. Notice that what I'm really saying here is that alpha 1 is 1u1 plus 1u2 plus 2u3 plus 3u4, alpha 2 is 2u1 plus 3u2 plus 4u3 plus 5u4. In essence, what I'm saying then is what?

Once I've specified a particular basis, every 4-tuple notation is relative to that basis. And by the way, what do I mean by being vague over here when I say, describe the space? What I mean is, for example, is the subspace all of v?

In other words, we know that the dimension of v is four. I have four vectors. If those four vectors are linearly independent, they will span all of v. I guess an equivalent question would be, therefore, are the alphas linearly independent?

I can make this even more vague. I can say, if w doesn't equal v, what does w equal? In other words, what subspace is spanned by the vectors alpha 1, alpha 2, alpha 3, and alpha 4?

Now, what I'd like to do is to introduce to you at this time a rather nice computational technique. It's nice for two reasons. The first reason is that it gets the job done very nicely, which is perhaps the best reason of all. The second reason, which is also very nice, is the fact that we've already had the technique before.

Keep in mind the following property of spanning vectors. If we have a set of vectors and we talk about the space spanned by those vectors, notice from our previous homework exercises that we have shown, that if you replace any member of that spanning set by itself plus a multiple of any other member of the spanning set, you do not change the space spanned by those vectors. Oh, again, by means of an example-- because as I say, this sounds very difficult from an abstract point of view-- but if you use three-dimensional space as an analogy, as a reference point, I think you'll see more clearly what these things mean. For example, i, j, and k are a basis for three-dimensional space. i, j, and k span three-dimensional space.

Suppose I replace i, or for the sake of argument, by i plus 3j. Look at the three vectors-- i plus 3j, j, and k. Those three vectors span the same space as i, j, and k, namely, all of three-dimensional space.

In other words, i plus 3j and j are non-parallel vectors in the x, y plane. Consequently, they spend the entire x, y plane. And then with the k vector, all of three space is spanned. So the idea is, that if ever we replace a vector by itself plus a multiple of any other factor in the spanning space, we do not change the space spanned by the vectors.

And my claim is that this suggests the row reduced matrix technique. Now, why is that? Well, the best way to see this is let's write, as a coding system, the alphas as a matrix.

In other words, what we'll do is is we will label the columns of our matrix, u1, u2, u3, and then the rows of our matrix will simply be 1, 1, 2, 3, 2, 3, 4, 5, 3, 7, 6, 5, 4, 5, 9, 9. In other words, in terms of our code, we have what? Alpha 1 is 1u1 plus 1u2 plus 2u3 plus u4, et cetera. So this is my coding matrix.

Now, what did we do when we row reduced? For example, to get a 0 in this entry-- remember our old technique? We said something like, let's take the second row and replace it by the second minus twice the first.

In terms of our coding system, this says, let's take the spanning vectors, alpha 1, alpha 2, alpha 3, alpha 4, and let's replace alpha 2 by alpha 2 minus twice alpha 1. You see, we are replacing alpha 2 by alpha 2 plus a scalar multiple-- the scalar happens to be minus 2, but that's irrelevant-- plus a scalar multiple of alpha 1. So the space spanned will not have changed.

In other words, if I now go through this row reduction, what do I do? I get 0s every place here. How do I do that?

I replace the second row by the second minus twice the first, the third row by the third minus 3 times the first, the fourth row by the fourth minus 4 times the first. In terms of the coding system, what am I doing? I'm replacing alpha 2 by alpha 2 minus twice alpha 1. I'm replacing alpha 3 by alpha 3 minus 3 times alpha 1. I'm replacing alpha 4 by alpha 4 minus 4 times alpha 1.

When I do that, my resulting matrix looks like this. See, notice now what's happened is these columns are still u1, u2, u3, and u4. But notice now, that whereas this vector here is still alpha 1, the vector 0, 1, 0, minus 1, in other words, what, u2 minus u4, is no longer alpha 2.

In fact, what is it? How do we get this row? It was alpha 2 minus twice alpha 1.

In a similar way, this would be what? 0, 4, 0, minus 4. But it's also what? Alpha 3 minus 3 alpha 1, and the vector 0, 1, 1, minus 3 is alpha 4 minus 4 alpha 4.

What we're saying is these four vectors may not look the same as our original vectors, but the thing that we're sure of is that they span the same space. Now, at this point, let me pause for a moment and try to give you some pseudo motivation as to why I'm doing this. You see, in the same way that we solved systems of linear equations, when we said that the same system of equations could be represented in different ways, in other words, two different systems could have the same solution set, but one of the systems was easier to solve than the other. See, one thing I'm saying over here is that, somehow or other, these four vectors look perhaps a bit simpler than these four because of the fact that I have 0s appearing over here. Well, I'll go into that in more detail later.

The other thing is this. Suppose, by row reducing this way, suppose I were to eventually wind up with the 4 by 4 identity matrix. In other words, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1-- what would that matrix represent? Notice, in terms of this coding system that 1, 0, 0, 0 would just be u1. 0, 1, 0, 0 would just be u2, et cetera.

In other words, if by row reducing this I could ultimately wind up with the identity matrix, that would say that the space spanned by alpha 1, alpha 2, alpha 3, and alpha 4 is the same as the space spanned by u1, u2, u3, and u4. But by definition of the us, those four vectors span my entire four-dimensional space. In other words, if in row reducing my alphas I wind up with the identity matrix, that will be proof, in terms of my coding system, that the alphas span all of v, because they will span the same space as u1, u2, u3, and u4. So maybe that's a more immediate reason for proceeding the way we're going.

At any rate, what I'm going to do now is to continue on with my row reduction scheme. In other words, what will I do now in terms of reduction? I will replace the first vector here by the first minus the second.

In other words, I'm replacing alpha 1 by alpha 1 minus this vector, which is alpha 2 minus 2 alpha 1. The important point being what? I'm replacing a spanning vector by itself plus a scalar multiple of another one.

In any event, going through this operation what I will end up with now is the following 4 by 4 matrix. And you see, again, what I'm saying is that the 4-tuples 1, 0, 2, 4, 0, 1, 0, minus 1, 0, 0, 0, 0, and 0, 0, 1, minus 2 span the same space as alpha 1, alpha 2, alpha 3, alpha 4. Now let me keep emphasizing that. Even though I said a 4-tuple always means with respect to u1, u2, u3, u4, unless otherwise specified, I think I will always specify it so we make sure that we get this idea hammered home.

Now, here's the first key point. My dream of having this matrix row reduced to the identity matrix is now shattered. Why is it shattered? The third row is already all 0s, so I can't possibly get a 1 into this position here. In other words, I can no longer hope to get the identity matrix here.

But this tells me something much more than that. In other words, even failure is success in this situation. Namely, the fact that this is a 0 vector tells me that the space-- the space spanned by alpha 1, alpha 2, alpha 3, and alpha 4 is certainly the same as the space spanned by these four vectors. But one of these vectors is the 0 vector, and the 0 vector spans no space, as we've already seen.

In other words, the 0 vector is always redundant, because any scalar multiple of the 0 vector is still the 0 vector. So that doesn't change what any other linear combination would be. So I immediately know now that the space spanned by alpha 1, alpha 2, and alpha 3 can't be any more than three-dimensional space. In other words, it's also spanned by these three vectors.

By the way, let me just continue to row reduce this for the sake of argument anyway. In other words, let me now replace the first row by the first minus twice the fourth. Or if I delete this, twice the third, but that's not important. If I do this, notice that the matrix I get now is 1, 0, 0, 8, 0, 1, 0, minus 1, a 0 row, and then 0, 0, 1, minus 2.

Let me give the three non-zero vectors here special names. Let me call this vector beta 1, this vector, beta 2, this vector, beta 3. And what does that mean? It means what?

That beta 1 is u1 plus 8, u4. Beta 2 is u2 minus u4. Beta 3 is u3 minus 2u, 4. In other words, this is beta 1, beta 2, and beta 3.

What do we know about beta 1, beta 2, and beta 3? Because they were obtained from the matrix of the alphas by row reduction, they span the same space as alpha 1, alpha 2, alpha 3, alpha 4. In other words, the first thing we know is that the space spanned by alpha 1, alpha 2, alpha 3, and alpha 4 is not four-dimensional space, but rather, it's the space spanned by beta 1, beta 2, and beta 3.

By the way, that immediately tells me, since beta 1, beta 2, and beta 3 span my space, w, that I'm looking for, that the dimension of that space can be no greater than 3, because 3 vectors span it. Now, at this stage of the game, the question might come up, what's so great about the betas? Why do we have the betas rather than the alphas? Do the betas help us at all?

And my claim is, making this metaphysical for a moment, if you just look at what the betas look like over here, just look at what the betas look like, forget about the u4 component, notice that the betas without the u4 component look like the identity matrix. See, 1, 0, 0, 0, 1, 0, 0, 0, 1. My claim is that the betas are, indeed, a very special basis. And let me show you what I mean by that.

What does it mean for a vector to be in the space spanned by beta 1, beta 2, and beta 3? Well it means that that vector must be a linear combination of beta 1, beta 2, and beta 3. Let me write down that linear combination in the form x1 beta 1 plus x2 beta 2 plus x3 beta 3, where x1, x2, x3 represent arbitrary constants here. By the way, don't be upset that I've used xs now for the first time instead of cs. I've essentially done the same thing that we did when we were dealing with arrows in the plane.

When we were talking about a typical arrow, we refer to it as being xi plus yj rather than refer to it as being ai plus bj to indicate, somehow how or other, that we had arbitrarily chosen the constants, the coefficients. It's sort of like in algebra. There's no law that says, when you have a constant, that you couldn't call the constant x, or if you have a variable, you couldn't call the variable c.

But somehow or other, when we see a c, we assume it's a constant in algebra. When we see an x, we assume it's a variable. The xs are just thrown in here to suggest that I'm looking at all linear combinations of beta 1, beta 2, and beta 3.

And now, watch how beautifully the form of the betas comes to help us now. Let's look explicitly to see what x1 beta 1 plus x2 beta 2 plus x3 beta 3 really is. Let's come back to what our betas were.

To multiply beta 1 by x1 means that we multiply each component by x1. So this will be what? x1, 0, 0, 8x1. We multiply each component of beta 2 by x2.

That's 0, x2, 0, minus x2. We multiply each component of beta 3 by x3. So I'll get 0, 0, x3 minus 2x3.

Notice, again, that these components as 4-tuples are relative to the us. In other words, x1 beta 1 plus x2 beta 2 plus x3 beta 3 are the three 4-tuples given here, where it's understood that the columns values are u1, u2, u3, and u4. Now, I want to add these up.

If I add them up, I add component by component. Here's the beauty of that row reduced matrix form that yielded the betas. Notice, that when I add these up, x1 appears only in the first vector. There are 0s every place else, x2 in the second, x3 only in the third.

Notice, that when I add these up, I get what? x1, x2, x3, 8x1 minus x2 minus 2x3. In other words, any linear combination of the betas in terms of the us has a very interesting representation. Namely, x1 beta 1 plus x2 beta 2 plus x3 beta 3 is simply x1, u1 plus x2, u2 plus x3, u3. Then plus 8x1 minus x2 minus 2x3, x4.

By the way, as an aside, notice that I can now very quickly show you that the betas are linearly independent. Namely, for linear independence, notice that I want to show that the only way this can be 0 is if all the xs are 0. But notice that since this combination comes out to be this, the only way this can be 0 is to have 0s every place. That says 0 must equal x1, 0 equals x2, 0 equals x3. In other words, writing this out in more detail, x1 beta 1 plus x2 beta 2 plus x3 beta 3 equals 0 means this, which in turn means that x1, x2, and x3 must all be 0.

That proves, in particular, that the betas are linearly independent. And therefore, my space, w, not only is spanned by beta 1, beta 2, and beta 3, but they are a basis. That's why I write the square brackets here. And so my dimensions of w is actually 3.

And by the way, let me show you, again, what this thing means. It means that, you see, if something isn't w, If it's written as a 4-tuple, all I have to do is know what the first three components are, because the fourth one is expressed in terms of the other three. For example, the u1 cannot belong to my space, w. Why? Because u1 has x1 equal to 1, x2 equal to 0, x3 equal to 0.

And this tells us what? That when x1 is 1, and x2 and x3 are 0, x4 must be 8. In other words, the only member of w that begins with 1, 0, 0, is 1, 0, 0, 8, and that's exactly what beta 1 is. In other words, I leave it for you the check that u1, u2, and u3 do not belong to us space, w.

Well, so much for that. Let me review now very quickly what we've shown. These betas are a rather remarkable basis. In other words, I claim that beta 1, beta 2, beta 3 is a natural basis for w, you see, much nicer than the alphas. Why is that?

Well, let me pick any 4-tuple in w. Let me call it x1, x2, x3, x4. My claim is that once I know that that 4-tuple is in w, I can write down the linear combination of the betas that this vector is by inspection. Namely, what I just showed was that that 4-tuple is what? It's x1 beta 1 plus x2 beta 2 plus x3 beta 3.

Or equivalently, if you want to see it in terms of the 4-tuple of notation in terms of the us, the first three components can be chosen at random, and then the fourth component is given by this. And by way of a quick example, notice that as a 4-tuple alpha 1 was what? It was 1, 1, 2, 3. In other words, it was u1 plus u2 plus 2u3 plus 3u4.

Notice that because the first three coefficients were 1, 1, 2 relative to the betas, the coefficients are also 1, 1, 2. In other words, I can write alpha 1 as the 3-tuple 1, 1, 2, provided I tacitly understand here that I'm making reference to the special betas as being my basis. You see, in other words, referring to the alpha 1, alpha 2, alpha 3, alpha 4 that we're talking about, just drop the fourth-- in terms of the 4-tuples, drop the use of four coefficient and just write down the remaining 3-tuple.

In other words, alpha 2 is 2 beta 1 plus 3 beta 2 plus 4 beta 3. Alpha 3 is 3 beta 1 plus 7 beta 2 plus 6 beta 3. Alpha 4 is 4 beta 1 plus 5 beta 2 plus 9 beta 3.

You see, the betas are a very nice basis, because they're that special basis that once I know what the vector looks like as a 4-tuple with respect to the us, I can immediately write down what it looks like with respect to the betas. Let me give you another example. Suppose I just pick at random-- well, it's not picked at random. I picked one that works here.

But let's suppose I just take the 4-tuple 2, 5, 3, 5, and I want to know whether that's in w. And by the way, again, let me remind you just one more time that 4-tuple stands for what? 2u1 plus 5u2 plus 3u3 plus 5u4. Does that 4-tuple belong to w?

The answer is, it belongs to w if and only if it's equal to what? 2 beta 1 plus 5 beta 2 plus 3 beta 3. That's the beauty of that special basis of the betas. Well, 2 beta 1 is simply 2, 0, 0, 16. 5 beta 2 is 0, 5, 0, minus 5, and 3 beta 3 is 0, 0, 3, minus 6.

Recalling, of course, that beta 1 was 1, 0, 0, 8. Beta 2 was 0, 1, 0, minus 1, and that beta 3 was 0, 0, 1, minus 2. So again, this is all with respect-- as 4-tuples with respect to the us.

If I add these up, I get what? 2, 5, 3, 5, which is exactly the vector I'm dealing with. So 2, 5, 3, 5 belongs to w.

In fact, another way of looking at this is once the 2, 5, and the 3 are given as the first three components, the only way that the 4-tuple can belong to w is if that fourth component is 5. In other words, again, looking at this from one more point of view, remember over here we showed that the fourth component in terms of the us had to be eight times the first minus the second minus twice the third. And given 2, 5, 3, the only way that can happen is what? 8 times 2 minus 5 minus twice three and that happens to be five.

Now we come to the home stretch. And that is, knowing now what the basis looks like in terms of the betas, we come back to a more crucial question. And that is, look it, we were given alpha 1, alpha 2, alpha 3, alpha 4. Those were the vectors that we were given. Maybe for the problem that we're dealing with, it's important to have everything relative to the alphas.

So now what we're saying is, knowing what 2, 5, 3, 5 looks like as a linear combination of the betas, how can we conveniently express that as a linear combination of the alphas? And again, the answer comes from row reduced matrices. In other words, all we're going to say is this.

Look it, we know what this vector looks like in terms of the betas. It's 2 beta 1 plus 5 beta 2 plus 3 beta 3. Suppose we could now express the betas in terms of the alphas. See, sort of like before with the matrices. We sort of want to talk about inverting-- interchanging the roles of the vectors.

Here's what I'm saying. Look it, we already know from over here that alpha 1 is beta 1 plus beta 2 plus 2 beta 3, et cetera. Let me do the following thing. Let me write down an augmented matrix.

The first three columns will stand for beta 1, beta 2, and beta 3, and the last four columns will stand for alpha 1, alpha 2, alpha 3, and alpha 4. Notice that I can write alpha 1 by two different names. It's 1, 1, 2 relative to the betas and 1, 0, 0, 0 relative to the alphas. And similarly, alpha 2, alpha 3, alpha 4 can be represented as I've written these things over here.

Now, what have I done? I have written each vector in two different ways-- one relative to the betas and one relative to the alphas. If I now row reduce this matrix-- notice that I'm dealing with vectors-- whatever vector I get by one name must be the same vector with respect to the other name.

In other words, if I start row reducing now in the usual way, this matrix becomes this matrix. And again, to show what we're saying is, this is just another way of saying what? That, for example, this is beta 2, and this says that the vector beta 2 is the same as alpha 2 minus 2 alpha 1.

This happens to be-- oh, let's look at this one. This is beta 2 plus beta 3. And that's the same as what? Minus 4 alpha 1 plus alpha 4.

And you see, what I'm going to try to do is row reduce this part here and hopefully be able to express the betas in terms of the alphas. So what I now do is continue my row reduction technique. And you say, if I do that, my next step gives me this, and this is a rather remarkable thing.

Because you'll now notice that reading my first three components here, this tells me that the third vector, at least in terms of the betas, is the 0 vector. By the way, what's another name-- see, what I'm saying this is what? 0 beta 1 plus 0 beta 2 plus 0 beta 3. What's another name for that vector? It's 5 alpha 1 minus 4 alpha 2 plus alpha 3.

In other words, as an aside, 5 alpha 1 minus 4 alpha 2 plus alpha 3 is the 0 vector. Therefore, by transposing, alpha 3 is 4 alpha 2 minus 5 alpha 1. And now, indeed, we have not only seen that the alphas were linearly dependent, we have found the redundancy. That in the order given, alpha 3 is a linear combination of alpha 1 and alpha 2.

By the way, just as a computational check, notice by our definitions of alpha 1 alpha 2 that 4 alpha 2 is 8, 12, 16, 20. Minus 5 alpha 1 is minus 5, minus 5, minus 10, minus 15. If I add these two, I get 3, 7, 6, 5, which, as you will recall from looking at our definition of alpha 3, is, indeed, alpha 3.

At any rate, what I'd not like to point out is that what this piece of information tells me is that this is the redundant piece of information. That what we really know is that the betas do not need alpha 3 to express them. Why? Because the betas could originally be expressed in terms of alpha 1, alpha 2, alpha 3, and alpha 4.

So you don't lose track of that. The space spanned by the betas is the same as the space spanned by the alphas. So, in particular, each of the betas must be a linear combination of the alphas by definition of spanning space, just as each of the alphas must be a linear combination of the betas.

So I know that the base can be expressed in terms of the alphas. But now, I've seen that alpha 3 is redundant, because alpha 3 can be expressed in terms of alpha 1 and alpha 2. So what I now have is what? That beta 1, beta 2, and beta 3 can be expressed in terms of alpha 1, alpha 2, and alpha 4.

See, notice that with this row deleted, the alpha 3 column has all 0s in it. And by the way, if I now finally complete my row reduction, you see, I wind up with this matrix, this 3 by 7 matrix. And this tells me, in particular, what? How to express beta 1, beta 2, and beta 3 in terms of the alphas, in particular, beta 1 is the 4-tuple 7, 1, 0, minus 2.

And notice now, I don't want to say this as a 4-tuple in the following sense. Notice, if I now were to write 7, 1, 0, minus 2, that would be redundant. Because now I'm talking relative to the alphas, not the us. The important thing is what?

That beta 1 is 7 alpha 1 plus alpha 2 minus 2 alpha 4. Beta 2 is minus 2 alpha 1 plus alpha 2-- well, just that's all it is. And beta 3 is minus 2 alpha 1 minus alpha 2 plus alpha 4. This tells me that the betas are a basis, the alphas weren't the basis, but alpha 1, alpha 2, and alpha 4 as a basis.

Getting back to our original problem, now that we're expressed beta 1, beta 2, and beta 3 in terms of alpha 1, alpha 2, and alpha 4, knowing that 2, 5, 3, 5 is 2 beta 1 plus 5 beta 2 plus 3 beta 3, we can now replace beta 1 by what it's equal to in terms of the alphas, beta 2 by what it's equal to in terms of the alphas, beta 3-- well, beta 2 is equal to minus 2 alpha 1 plus alpha 2. Beta 3 is minus 2 alpha 1 minus alpha 2 plus alpha 4.

Replacing the betas by what they're equal to in terms of the alphas, I wind up with this expression. And I find that 2, 5, 3, 5 is the following linear combination of alpha 1, alpha 2, and alpha 4. Namely, minus 2 alpha 1 plus 4 alpha 2 minus alpha 4. And I leave it for you to check that, if you actually take alpha 1, alpha 2, and alpha 4 as given at the beginning of this example, you will indeed find that this linear combination, minus 2 alpha 1 plus 4 alpha 2 minus alpha 4, is indeed the vector 2, 5, 3, 5 that we're talking about.

Well, as I mentioned to you, I want you to get an overview of what we're talking about. There are going to be plenty of exercises on this. Let me just summarize, or at least give a partial summary, of what we have really done computationally in this example. We started with four-dimensional space, v, relative to a basis u1, u2, u3, u4.

We picked four vectors-- alpha 1, alpha 2, alpha 3, alpha 4-- in that space, and asked to see what space was spanned by those four vectors. And we showed that there were three very special vectors-- they had a very nice form-- called beta 1, beta 2, beta 3, such that v had-- the w had beta 1, beta 2, and beta 3 as a basis so that the dimension w was 3. We also showed in this particular problem that not only was the dimension three, but once you picked the u1, u2, and u3 components arbitrarily, the u4 component was determined by these in this particular way. In other words, another way of saying this is that w consisted of all linear combinations of the betas, where these are the same xs that appear here.

And as a case in point, we showed that a particular vector, 2, 5, 3, 5, belonged to w. That when we wrote that vector relative to the betas, that vector was 2, 5, 3. When we wrote that same vector relative to the alphas, it was minus 2, 4, minus 1.

Notice, by the way, a very important thing here. You see, every single time that we talked about two vectors are equal if and only if they were equal component by component, when we wrote them as n-tuples, it was assumed that the basis never changed. Look it, forget about the betas and the alphas here. Look at the 3-tuple 2, 5, 3 and the 3-tuple minus 2, 4, minus 1. Certainly, they are different 3-tuples, but they happen to name the same vector, because one is a 3-tuple coded relative to the betas and the other is a 3-tuple coded relative to the alphas.

In other words, once we pick a particular basis, we must remember that our coding system is relative to that basis. If, for some reason, it becomes convenient to change the basis, then we have to be very, very careful and remember that, therefore, our coding system has changed. And that's why, in most advanced applications using vector spaces, a great degree of difficulty comes in, because what we've done is we've changed basis vectors. And people who are so rigid that they keep visualizing that you are using the same n-tuples over and over again get confused as to how two different n-tuples can name the same vector.

Well, at any rate, I think that's more than enough for one session. Do the exercises very carefully. Try your best.

As I say, if you have any trouble after the exercises, review the film again. Listen to the lecture again. And what we will do next time is pick up some consequences of vector spaces that tie-in with other aspects of the course that we've already studied. But at any rate, 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.