HERBERT GROSS: Hi. As I was getting myself prepared for the lecture, an old shaggy dog science story came to mind that's usually told as a tribute to the German scientist's thoroughness. This was the story of when all the scientists of the world got together and decided for an annual project. Each country would study exhaustively a different animal and report at the end of the year as to what they had found out.

And at the end of the year, every country but Germany had been heard from. No one knew what had happened to the German scientists. And five years later, the German report came in, a huge epitome, and it was entitled, "Handbook of the Elephant, Volume 1." And the reason that this story comes to mind is twofold.

First of all, the study of vector spaces has so many ramifications, that to study the subject thoroughly should be a full year's course by itself. And secondly, with respect to the particular topic of basis vectors and spanning vectors, and whether subelements are linearly independent, whereas we've invented a rather nice row reduced matrix technique for finding out what space is spanned by a set of vectors and what a basis for that space is, frequently we don't want that much information. And what I'm leading up to is a topic which you've all had in a previous context.

It's called determinants. And today, what I would like to do is to study determinants within the framework of vector spaces. And the way it comes up is as follows.

Suppose we have an n-dimensional vector space, v, and suppose that we pick a particular basis, u1 up to un, for v itself. Now, the idea is this. Knowing that the dimension of v is n, we immediately know that more than n vectors can't be a basis for v, because they will be linearly dependent rather than linearly independent.

Fewer than n vectors of v cannot be a basis, because fewer than n vectors, since the dimension is n, cannot span v. Consequently, the only point of interest is what happens when we're given a set of n vectors and v. Namely, given the n vectors, alpha 1 up to alpha n and v, is this set a basis or isn't it?

And rather than to use the row reduced matrix technique, here what we're saying is we don't want, for example, in many cases, to know what the betas look like that we talked about in our lecture on spanning vectors and the like. All we want to know is, are these vectors linearly independent or aren't they?

And so what we do is we essentially invent a function machine. In other words, what we're going to do is to construct a function in which, given that the dimension of our space is n-dimensional, the input of our function machine will be any set of n vectors from that space. And the machine will be programmed to give 0 as an output if the vectors are linearly dependent, in other words, if they are not a basis, and non-zero-- and I'll explain in more detail what non-zero-- why I picked non-zero rather than a specific non-zero number-- if the n vectors do form a basis.

In other words, I am going to invent a function, D, capital D, to indicate the word determinant here, such that D of the n vectors, alpha 1 up to alpha n, will be 0 if and only if the set of n vectors are linearly dependent. Now, obviously, a function machine can do nothing by itself. We have to give-- and I hope this is well-known by now in our course-- we have to endow anything that we're working with with a particular structure so that is feels free to work logically for us. Well, among other things, what do we know for sure is a set of n vectors which are a basis for v?

Since we were given that the specific basis that v is being referred to with respect to are u1 up tp un, then certainly, if u1 up to un is the input of the D machine, we want a non-zero output. For the sake of just normalizing things, let's program the machine so that D of u1 up to un will be 1. Now, notice, by the way, this only tells us one special basis. There are many bases that I could have chosen for v, and certainly condition one by itself isn't going to tell me any-- has no information programmed in it to tell me anything other than what it does to u1 up to un.

As a second input to my D machine for programming it, I certainly know that given the n vectors, if any two of the vectors happen to be equal, then those vectors are linearly dependent. Consequently, I instruct my D machine, that if the input is a set of vectors alpha 1 up to alpha n, and at least two of the alphas are equal-- and the way we say that mathematically is that alpha i equals alpha j for a sum i unequal to j-- that if two of the vectors in the set are equal, we tell the machine to grind out 0 as an output. And by the way, at this stage we should be very careful to recognize that this is not the only way in which a set of vectors can be linearly dependent. There are many ways in which a set of vectors can be linearly dependent, even if no two of the vectors are equal.

So consequently, given a linearly dependent set of vectors in which no two are equal, notice that condition two here will not work or do anything for me. So all I have done by conditions one and two is I have endowed my D machine with two standards. One particular case in which the output of the D machine will be the number 1, and a particular situation in which the output of the D machine will be 0. But I have not solved the problem that I wanted so far, namely, to make sure that D grinds out 0 if and only if the set of n vectors are linearly dependent.

Let me now endow D with a third property. It is a property that will suggest linearity to you, and it's one which is amazingly powerful. And what I mean by amazingly powerful hopefully will become clear in a few moments. But the third property is this.

And rather than work with n-tuples here, to make things easier to read, let me just take a 3-tuple over here. Suppose I have a three-dimensional vector space, and I pick three vectors in that space. Suppose one of those vectors-- the way I've written that here, it happens to be the second vector. Suppose one of those vectors is itself written as a sum of two other vectors in the space.

Then I tell the D machine to compute this as follows. Linearize this, in other words, compute this as if it were the sum of two separate determinants, one of which had the alpha 2 missing, and the other of which had the beta 2 missing. In other words, what I do is is I endow the machine with the property that D of alpha 1 comma, alpha 2 plus beta 2 comma, alpha 3 will be D of alpha 1, alpha 2, alpha 3 plus D of alpha 1, beta 2, alpha 3.

And thirdly-- well, thirdly. I don't mean thirdly. I mean the other lineal property. If one of the input vectors is multiplied by a scalar, I can factor the scalar out.

In other words, if the three vectors that are being tested in my three-dimensional space have the form alpha 1, c alpha 2, and alpha 3, D of alpha 1 comma, c alpha 2 comma, alpha 3 will be-- see, factor of the c out. c times D of alpha 1, alpha 2, alpha 3. And it's extremely important to notice that the c does not have to be a scalar multiple of all three vectors in the input.

Notice, that the fact that the c was multiplying one of the vectors is enough to factor the c out. Without belaboring this point, all I wanted to say was, that if the c were multiplying each of the alphas, in other words, if this were D of c alpha 1 comma, c alpha 2 comma, c alpha 3, this would equal c cubed D alpha 1, alpha 2, alpha 3, because we would fact about the c each time-- one for each vector. But let me illustrate this in terms of a 2 by 2-- a two-dimensional example for you.

Suppose I'm dealing in two-dimensional space, and I'm dealing with respect to a particular basis, u1 and u2. And let's suppose that relative to that basis, u1 and u2, alpha 1 is the vector 3 comma, 1, beta 1 is the vector 6 comma, 7, and beta 2 is the vector 4 common, 5. Let's recall the fact that, one way or another, we already know how to expand 2 by 2 determinants, even though we may not know rigorously why the rules were chosen.

The idea is, what would alpha 1 plus beta 1 be? It would be the 2-tuple whose first entry was 3 plus 6 and whose second entry was 1 plus 7. In other words, the first row of the determinant that I'm going to be talking about has as its entries 3 plus 6 and 1 plus 7. Beta 2, which will make up the second row with my determinant, has as its components 4 and 5. So the second row of my determinant is 4, 5.

And using the usual rule for multiplying determinants, this is 5 times 9 minus 4 times 8, which is 13. And by the way, just to refresh your memories here, what if I computed in terms of my D language and the alpha 1, alpha 2, and beta 2 is used here? This is D of alpha 1 plus beta 1 comma, beta 2.

On the other hand, what would D of alpha 1, beta 2 be? D of alpha 1 beta 2 would be the 2 by 2 determinant, whose first row was 3, 1, and whose second row was 4, 5. D of alpha 1, beta 2 would be the determinant whose first row-- I'm sorry. D of beta 1 comma, beta 2, would be the determinant whose first row was 6, 7 and whose second row was 4, 5.

In other words, these matrices here, what is this again? This is D of alpha 1, beta 2. And this one here is D of beta 1, beta 2.

And notice, that this determinant by the traditional way of expanding is 11. This determinant is 30 minus 28, which is 2. 11 plus 2 is 13. And we see that, at least in this case, D of alpha 1 plus beta 1 comma, beta 2, is the same as D of alpha 1 comma, beta 2 plus D of beta 1 comma, beta 2.

So at least the first part of property three is obeyed in this particular example. And to show what the second property is, let's suppose alpha 1 now-- I pick a new alpha 1. We'll call that 2 comma, 6. Suppose alpha 2 is 3 comma, 4.

Then the determinant of alpha 1 and alpha 2 would be what? 2, 6, 3, 4. And that's what? 8 minus 18, which is minus 10.

On the other hand, notice that a common-- that this vector here could have been written as twice 1 comma, 3. I can, therefore, factor out the 2. See, notice this very carefully.

Notice that the alphas, which are written in a row here-- one row this way-- each alpha forms a row when I use the matrix interpretation. So what I'm saying is, that 2 comma, 6 is one vector. It's a 2-tuple with respect to the basis u1 and u2. All I'm saying is factor out the 2 from the first row, from the first vector.

That leaves me with 2 outside, 1, 3, 3, 4 inside. And notice, that 4 times-- 4 minus 9 is minus 5 times 2 is also minus 10. Notice that these, indeed, are equal, and notice that I could factor out the 2 simply by virtue of the fact that it was a common factor in one row.

It did not have to be a common factor in both rows. In fact, if it had been a common factor in both rows, I would have had to factor out a 2 twice, in other words, a 4 over here. Now, I can go on with things like this, but, again, I want to stress the overview.

And the key point is that these three simple properties, which I've just called simply one, two, and three, those three properties that I've programmed the D machine with are enough to completely determine D. In other words, if I wanted to be dramatic over here, the key point is that one, two, and three completely determine D. I'll put an exclamation point down there.

Because what I claim is, that once these three properties are obeyed, there is no possible way for D to behave other than in a very unique well-defined way. That D now has a perfectly well-defined structure. Let me give you some examples of that structure.

I'll just prove a couple of theorems. I'll prove them in two-dimensional space, leaving it for the exercises to work on higher dimensional space, but this doesn't get too cluttered. First of all, what I claim is that the D machine is so finicky, that if you interchange the order in which the two vectors are given, you change the sign of the output. In other words, D of alpha 1, alpha 2, in that order, is the negative of D of alpha 2, alpha 1.

And the proof, again, follows very nicely in terms of our game structure. The gimmick to begin with is that we must be clever enough to decide that we'll compute the determinant of alpha 1 plus alpha 2 comma, alpha 1 plus alpha 2. On the one hand, by our second property, since two of the vectors making up the set being tested are equal, it means that that determinant must be 0.

On the other hand, by property three, by linearity, splitting this up as two terms, notice that this is D of alpha 1 comma, alpha 1 plus alpha 2 plus D of alpha 2 comma, alpha 1 plus alpha 2. In other words, this result here. Now, in turn, noticing that these form a sum, I can rewrite this one as what? D of alpha 1 comma, alpha 1 plus D of alpha 1 comma, alpha 2.

This one is D of alpha 2 comma, alpha 1 plus D of alpha 2 comma, alpha 2. I've just written this whole thing out on this next line. By property two, D of alpha 1 comma, alpha 1 must be 0. D of alpha 2 comma, alpha 2 must be 0, because after all, you see, two vectors are equal in this.

So what do I have left? Comparing this with this, I have the D of alpha 1 comma, alpha 2 plus D of alpha 2 comma, alpha 1 is 0. And that means that the number-- and keep that in mind, the determinant is a number. It maps the n vectors into a number, which is either 0 or non-zero. But the determinant of alpha 1 comma, alpha 2, therefore, must be the negative of the term of alpha 2, alpha 1, because their sum is 0.

The second theorem-- and this is the one that relates determinants to matrices, a very fantastic result, one that has tremendous practical application that will also talk about later in the lecture. And that is, that if I take any one of my input vectors and replace it by itself, plus a scalar multiple of another, I do not change the determinant. In other words, for example, if I replace alpha 1 by alpha 1 plus a scalar multiple of alpha 2, and I leave alpha 2 alone, the determinant of alpha 1 and alpha 2 will be the same as the determinant of alpha 1 plus c alpha 2 comma, alpha 2. And again, the proof is very easy, namely, look at the expression D of alpha 1 plus c alpha 2 comma, alpha 2. Use the linearity property that this is D of alpha 1 comma, alpha 2 plus D of c alpha 2 comma, alpha 2.

Then use the second part of the linearity property that the constant factor, c, can be taken outside here. So this becomes D of alpha 1, alpha 2 plus c times D of alpha 2 comma, alpha 2. And notice, that since alpha 2 is repeated here, this determinant is 0. And consequently, what we've proven is that this is equal to this.

Now, I've just proven those two theorems, because that's about all I need. Let me show you now, that if I had never seen that shortcut method of expanding a 2 by 2 determinant, as learned in high school, how these three properties uniquely determine what D has to mean. In other words, why I said that these three properties uniquely determine D.

As an example, let me take a two-dimensional vector space again. Let me pick u1 and u2 as a specific basis, and that's important. I've picked a specific basis, u1, u2.

Now, I pick any two vectors. Since v is a two-dimensional space, I pick any two vectors-- alpha 1 and alpha 2. And because of this notation, this means that alpha 1 and alpha 2 are both linear combinations of u1 and u2. Say alpha 1 is this, and alpha 2 is this.

Therefore, what is D of alpha 1 comma, alpha 2? Well, by direct substitution, D of alpha 1, alpha 2, just replace alpha 1 by what it's equal to here, alpha 2 by what it's equal to here. And we get that D of alpha 1, alpha 2 is this expression.

Now, we use the linearity property. We treat this as one number for the time being and split this up as a sum. In other words, it's going to be D of this term comma, this whole term plus D of this term comma, this whole term. In other words, if I do that, this breaks down to this.

Now, I notice that each of these two is a 2-tuple in which the second entry is a sum of two terms. So I now split this up into what? D of this comma, this plus d of this comma, this.

This term becomes D of this comma, this term plus D of this comma, this term. And I hope it doesn't sound boorish on my part to keep saying "D of this comma, this." I prefer to say that for you so that you can watch what I'm doing, and then have you just be able afterwards to read what these things mean.

Now, when I'm down to here, notice that by my second part of linearity I can factor out the a11 from here. I can factor out the a21, because it's a common factor of this term. In other words, I can write this term as a11, a21, D of u1 comma, u2.

And again, sparing you the details, notice I can factor out an a11, a22 from this term. I can factor out an a12, a21 from this term, an a12, a22 from this term. And that this now can be written in this particular way.

Now, what properties do I know that D is endowed with? I know, first of all, that whenever D operates on a set of vectors where at least two of them are equal, D must give 0 as an output. So D of u1 comma, u1 is 0 again.

u1 comma, u2 is the particular basis with respect to which I define v, you see. Therefore, by property one, D of u1 comma, u2 must be 1. D of u2 comma, u1 is just the permuted order of u1 and u2. Therefore, by our first theorem, that must be minus D of u1, u2.

Therefore, it must be minus 1. And finally, D of u2 comma, u2 is 0. And if I now collect everything that I have left here, what do I have? I have the D of alpha 1 comma, alpha 2 is equal to what?

It's equal to a11, a22, see, times 1, minus a12, a21. In other words, this is the determinant of alpha 1 comma, alpha 2. And let me, again, put an exclamation point here. Because if you now go back to the high school way of computing this determinant-- remember how we did it? What would we have obtained?

And, again, let me just come back to this board for a second here. Remember how you used the matrix of coefficients? So take the determinant, would be what? a11, a22 minus a12, a21.

So, again, notice two things that happen here. First of all, I get the same answer as I would have got the traditional way. And second, just as a minor aside that I'll emphasize more later in the lecture and also in the exercises, notice that you can begin to suspect that the actual value of the determinant of alpha 1 and alpha 2 should depend on what basis was chosen. Because relative to a different basis, notice that the coefficients might very well be different for alpha 1 and alpha 2.

But I don't want to mention that right now. I think, as I say, what the amazing result is, that what I meant by saying that these three properties completely determine D, is the fact that with these three properties it turns out that the high school definition was ironclad. Meaning, there was no other possible way to define what the determinant should be if you wanted these three properties to be obeyed.

By the way, I have some quick checks. Notice that relative to the basis u1, u2, u1 can be written this way, u2 can be written this way. We know that D of u1, u2 should be 1. And using the old fashioned way for checking this, what would the determinant be? It would be 1 times 1 minus 0 times 0, which is 1.

Secondly, we also know that property two should yield 0 as a determinant if these two vectors that were making up the input were equal. Let's call the vectors au1, bu2. So the 2 by 2 determinant I would get in this case is ab, ab. And if I expand that determinant, it's what? a times b minus a times b, which is 0.

And now, let me come back to that idea of what motivated the whole block of material here, the trouble with writing vectors as n-tuples. That there really is something that requires great care. In advanced applications, one must always be wary of this. I am not going to give you the advanced applications here. All I want you to do is to become prepared against the pitfalls.

And the thing is, that in dealing with determinants, as I said before, you must be very, very careful about what basis you're referring to for a given vector space. Obviously, the vectors do not depend on the basis, but their representation does. Let me give you a for instance.

Using v and u1 and u2 as an example number 1, suppose we let alpha 1 be the 2-tuple, 3 comma, 4. In other words, 3u1 plus 4u2. And let alpha 2 be 2 comma, 5.

Then by what we've just proven, the determinant of alpha 1 comma, alpha 2 must be the determinant of what matrix? The one whose first row is 3 comma, 4, and whose second row is 2 comma, 5. If I compute that determinant-- and by the way, notice, now that I know that the shortcut way has to be the right answer, I don't do this the long way anymore. I just say, OK, it's 3 times 5 minus 2 times 4 15 minus 8, which is 7. So the determinant of alpha 1, alpha 2 is 7.

So far so good. But let me point out the following thing. Notice, that alpha 1 and alpha 2 are also linearly independent. Namely, 3 common, 4 is not a scalar multiple of 2 comma, 5.

Consequently, since alpha 1 and alpha 2 our linearly independent, and since v is a two-dimensional space, it means that alpha 1 and alpha 2 are themselves a basis for v. In other words, v is equal to the space spanned by alpha 1 and alpha 2 as a basis. Notice, that relative to this new basis, alpha 1 comma, alpha 2, alpha 1 is 1 alpha 1 plus 0 alpha 2, alpha 2 is 0 alpha 1 plus alpha 2.

And therefore, using the traditional method of computing the determinant, notice that the determinant of alpha 1 comma, alpha 2 relative to the basis alpha 1 and alpha 2 would be what determinant? It would be to the determinant of the matrix whose first row was 1, 0, and whose second role was 0, 1. And that is 1.

This looks like a contradiction. You see, on the one hand, we have the D of alpha 1 and alpha 2 is 1. But down here, we just saw that D of alpha 1 and alpha 2 was 7. Which is correct?

Well, the answer is, they're both correct. That the first observation is, that if the determinant of alpha 1 and alpha 2 is not 0, the value depends on the particular basis. In other words, if the determinant is 0, it turns out that no matter what basis you use to represent the determinant of alpha 1 and alpha 2, you will get 0 once the terminator is 0 in one basis.

And if the determinant is not 0 with respect to a particular basis, it will be non-zero with respect to all bases. But what non-zero number it will be does depend on the basis. And that was why way back at the beginning we told the D machine simply to give a non-zero output if the input was a set of n linearly-independent vectors.

In summary, when one says that the determinant of alpha 1 and alpha 2 is 7, it's tacitly assumed that D is being referred to with respect to the particular basis u1, u2. On the other hand, when one says that D of alpha 1 and alpha 2 is 1, it's tacitly assumed that the basis that we're using to express v is alpha 1 and alpha 2. Again, I'll drill that more in the exercises.

Let me go on now to generalize what happens in n-dimensional space, keeping in mind the fact that when we generalize what happens in n-dimensional space, the actual proofs become messier, but the theory remains very much the same. I prefer to leave the messy details to the exercises, either optional or required, depending on how hard the exercises may be. But let me summarize what the results are for any n-dimensional space, keeping in mind we've proven the results rigorously for the case n equals 2.

The generalization is this. If I have an n-dimensional vector space, v, with respect to a particular basis, u1 up to un, and if alpha 1 up to alpha n are n vectors chosen from v so that they are linear combinations of the us, say, in the traditional way that we've written this all the time, it turns out the following. That the determinant of alpha 1 up to alpha n-- first of all, it's conventional and convenient to write the determinant as if it were a matrix, only replacing the square brackets by sort of absolute value signs-- vertical lines. What we do is, notice that alpha 1 is written as an n-tuple, a11 up to a1n.

In other words, the first rule of the determinant represents alpha 1 as an n-tuple vector relative to the basis u1 up to un. And a similar thing holds for alpha 2 through alpha n. And you may remember, I told you what the recipe was when we were dealing with-- I told you what the recipe was when we were dealing with cross products and the like. And the idea is simply this.

What you do is is you start in the upper left-hand corner writing a plus sign. Then you alternate going along rows and columns in any order you want writing plus, minus, plus, minus, et cetera, plus, minus, plus, minus, et cetera. Pick any row or column that you want. And then you go down that row or column factoring out a particular term, the plus sign telling you to take out the term as it is, the minus sign telling you to take out the term and change its sign, and you multiply that term by the n minus 1 by n minus 1 matrix that's left when you strike out the row and column in which that term occurs.

Now, if that still sounds like a tongue-twister, let's do this in terms of a specific semi-abstract, semi-concrete, four-dimensional case. Namely, let's suppose v is a four-dimensional vector space relative to a particular basis, u1, u2, u3, and u4. Suppose alpha 1, alpha 2, alpha 3, and alpha 4 are these four specific vectors of v. And again-- I can't keep emphasizing this too much, even though I hope some of you are bored because you know it so well-- it's always understood when I write this that the components are relative to u1, u2, u3, u4. For example, this is 3u1 plus 5u2 plus 7u3 plus 2u4.

Now, the traditional way of expanding this determinant is you write down the rows as the n-tuples. See, you write down-- this is your first row, second row, third row, fourth row. So you notice that the determinant is an n by an array of numbers that you have four vectors. Each vector is a 4-tuple, and this is how you get this determinant.

Now, what do we do next? Let's say, for the sake of argument, I elect to expand this determinant along the top row. I first take the 1 out. Because it has a plus sign, that comes out as 1.

I multiply that by the 3 by 3 determinant that's left when I strike out the row and column in which 1 appears. I take out 2, but make it a minus, because of the signature here, and multiply that by the 3 by 3 determinant that's left when I strike out the row and column in which 2 appears. I factor out 1 as it is, and multiply that by the 3 by 3 determinant that's left when I strike out the row and column in which this 1 appears. And finally, I take out this minus 1 and change its sign, because of this minus code over here. In other words, this comes out as a plus 1 multiplied by the 3 by 3 determinant that's left when I strike out the row and column in which minus 1 appears in.

And to make a long story short, all we're saying is that this 4 by 4 determinant can be written as a sum of four 3 by 3 determinants, namely, this, which is precisely what I was saying the long way up here. Now, what we say is, each of these 3 by 3s can be written as three 2 by 2s. Namely, I can factor out-- say again, I'm going across the top row.

I can pick any row or any column, but let me stick to the top row, like we did in i, j, and k, for the time being. But this becomes what? 5 times this 2 by 2 matrix, minus 3 times this 2 by 2 matrix, plus 1 times this 2 by 2 matrix. In other words, this expression here.

I can now do that for each of these. So each of the 3 by 3s gives me three 2 by 2s. So altogether, I would have 12 2 by 2s and go through this whole mess to compute this thing. And this can be very, very difficult, sure.

This was only a four-dimensional case. Imagine trying to apply this technique for a 10-dimensional vector space, for the sake of argument. It turns out that that recipe that we called theorem two gives us a tie-in between row reduced matrices and a quick way to compute determinants.

For example, let's suppose we still wanted to compute the same determinant that we've written down over here. Let's call that-- let's write it again over here. What we already know is, if I were to replace the second vector by the second minus twice the first, the determinant does not change. That's what theorem two said. If you replace one vector by itself plus a scalar multiple of another, you don't change the determinant.

Similarly, if I were to replace this vector by three times the first vector, I would still have the same determinant. And finally, if I were to replace this vector by three times the first vector, I would still have the same determinant. I don't know if you've noticed what I've been driving at. But wasn't what I was saying here the same computational steps that one goes through in row reducing a matrix?

In other words, let me row reduce this matrix so I get 0s every place, say, in my first column, except in the upper column. By the way, notice over here I always have the habit of putting a 1 in the upper left-hand corner. That's simply to facilitate the arithmetic. I hope by this stage of the game you realize, if this weren't a 1, I can always divide through by it, or what have you, and do the arithmetic another way.

I just do this thing for convenience. But the important point is that this determinant is now equal to this determinant. And be very careful.

Notice, if these had been matrices-- in other words, if I had put square brackets here and had made these matrices-- we would call the two matrices row equivalent but different matrices. Notice, the determinant is a number. And what we're saying is that the number named by this determinant is equal to the number named by this determinant-- equal.

Here's the key point. I now notice that I have three 0s in the first column. I, therefore, elect to expand this determinant along the first column. Why do I do this? Well, let's take a look.

First of all, I get a 1 multiplied by this 3 by 3 determinant that's left when I strike out the row and column in which 1 appears. The beauty is that when I form the other three determinants, 3 by 3 determinants, they're all going to be multiplied by 0. For example, over here, I factor out-- well minus 0, which is still 0.

What do I multiply that by? I multiply that by the 3 by 3 determinant that's left when I strike out the row and column that this 0 appears in, namely, the 3 by 3 determinant, 2, 1, minus 1, minus 1, 4, 5, 2, 3, 4. But since that's being multiplied by 0, the product will be 0.

To make a long story short, to evaluate this determinant, it's simply what? Plus 1 times this 3 by 3 determinant. I put the 1 in parentheses here to indicate that I really factored that 1 out. But really what we're saying is that this 4 by 4 determinant has the same value as this 3 by 3 determinant.

And now, I row reduce this 3 by 3 determinant. I replace the second row, in other words, the second vector, by the second plus the first, the third by the third minus twice the first. I now wind up with this 3 by 3 determinant.

Since I don't change the value of the determinant when I replace one vector by itself plus a scalar multiple of another, then notice very quickly that what happens over here is that these two determinants are equal. Again, I elect to expand along the first column factoring out the plus 1 being left with the determinant whose entries are 5, 8, 1, 1. And the other two terms contribute nothing, because the coefficient is 0.

But I already know how to expand-- I've proven that-- how to expand the 2 by 2 matrix, determinant. That's just going to be what? 5 minus 8 or minus 3. And that's how this shortcut row reduction works-- much, much more elegantly than the brute force technique. And by the way, this is the technique used in most, if not all, computers in determining the determinant of a set of n vectors in an n-dimensional vector space.

Well, as I say, this was meant as an overview. I hope you now see the overall picture of what determinants are all about. And in the exercises, I will try to give you some more sophisticated and elaborate details.

Next time what we will do is talk about an application of how determinants are used in certain aspects of vector spaces. In particular, we're going to talk about something called eigenvalues or eigenvectors, but more about that next time. 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.