[SQUEAKING]

[RUSTLING]

[CLICKING]

GILBERT STRANG: OK, this is a talk about one topic in the linear algebra book that people have appreciated. And that's the four fundamental subspaces. So in doing it, we'll understand better what a subspace is. And these four are associated with a matrix.

So we start with a matrix. And I want to describe its column space and its row space. Those may be not new to you. The null spaces of the matrix and its transpose, that superscript, t, means that the transpose exchanges the rows, makes them the columns, and the columns become the rows. So that's a transpose.

So we have two matrices. They each have a column space and a row space, so four fundamental subspaces, and then this great first fact about linear algebra, about the row rank equaling the column rank. So a matrix just has one rank.

Well, let's go, just to say where these topics apply the most. For square matrices, where you have maybe 12 equations and 12 unknowns, m and n, and the row count and the column count are the same, those are normal in physics and engineering. You're trying to find the stress on a structure. There, you have a square matrix that's usually invertible. Well, the four fundamental subspaces are still there, but two of them are empty, the null spaces. I'm thinking of this matrix as invertible, so the null space is only the 0 vector. And same for the transpose. So solving Ax equal b is a separate and fundamental problem.

Here are more statistics, like regression, like least squares, and data science now. So many applications. So I thought data science is so big that it would be worth focusing again on this case, where is A is a rectangular matrix. A rectangular matrix couldn't have an inverse on both sides because on one side, you put m by m, and on the other side, m by n. It can't be the same.

But there is something called the pseudoinverse. That's a little advanced, but you'll see what it's about. So we're in the data science world. Oh, and I thought I'd show a picture, the big picture, of linear algebra before. And then you'll see it again when you know more about it. But this is the picture that we're aiming to understand. The column space is up on the top right and the row space is on the top left. And again, for an invertible matrix, that's the whole space. That's the whole space.

But for data science matrices, we have null spaces. And those are the two with a capital N, for null space. And you'll see what those are. So that they complete the picture. So this is just a lovely picture of the subspaces associated with a matrix A. OK, so you'll see this again when we've gone into each of the four.

So here's the first, the column space of a matrix. So I have to take an example. So there's a matrix that has two rows and three columns. And we'll find the four subspaces for that matrix. OK. So in the first one is the column space. So what does the column space mean? So the columns are column vectors, 1,4, that's a vector in two-dimensional space. 2, 5 is a different vector. 3, 6 is another vector, all in two-dimensional space. And we take their combinations, all combinations, of v1 times the first column, v2 times the second, v3 times the third. That's a fundamental operation of linear algebra, multiplying a matrix by a vector.

And notice that I don't do it, the dot product, the standard way. I do it the linear combination way, writing this, keeping vectors in the picture, a combination of that. And we want to find, if I allow v1 and v2 and v3 to be any numbers, all numbers, for all v1, v2, v3 each choice gives us a point, and each which is in two dimensional space here, because these are vectors, or have two components.

And if we take all combinations, all the vs possible, that will fill up a plane. And in this case, there's only the xy-- All the vectors are in the xy plane. And those combinations will just fill the plane because 1, 4 is in an independent direction from 2,5. Actually, just the combinations of 1,4 and 2,5 will fill the plane. You've got to see that. You've got to picture that in your mind. The vector 1,4 points somewhere. The vector 2,5 points somewhere else. And the combinations of those two give all the vectors in between and all the negative vectors because the vs could be negative numbers.

So the first two already fill the whole xy plane. And so that's the column space for this matrix, is the whole plane. So that's an extreme case. We got the whole xy plane. And as it says at the bottom, the first two columns are already in different directions. Their combinations would produce the whole plane just with the first two. And the third is a little extra.

OK, good for the column space? So next space for this same matrix is the row space. But I like that letter C. And I don't want to introduce a brand-new letter if I don't have to. So what I do is transpose the matrix. "Transpose" means the matrix gets turned on its side. And 1,4 was the first column, and now it's the first row. And so now I have a different shape matrix. It has three rows and two columns. And now I'm taking combinations of those.

And what do I get? Here's a good example. I have two columns, 2 vectors, A1 and A2. And they are in three-dimensional space. So one vector in three dimensional space is just along a line. And if I took all the-- if I multiply it by anything, I'd get the whole line, but it would be just a line. But now, I throw in the second column. That, by itself, would give me a line in another direction. But when I take all points on one line plus all points on the other line, I fill a plane, a two-dimensional plane. I nearly said the whole space, but that's wrong.

So the row space here is a two-dimensional space, a plane inside 3D. All these columns have three components. So they're vectors in three-dimensional space. So we're in the space we live in, 3D. But the A1 and A2 and their combinations fill a plane in that space. So this is the second of our four subspaces. And you can read below, when I'm taking negative combinations, that gives me the vectors in the opposite direction. v1 and v2 can be any numbers, decimals, fractions, whatever.

So here's a picture of this space. I like this picture. I think it's got just about the right number of vectors to show you the combinations. This is only whatever it is, 25 combinations or something. No, just 20. So we have 20 points. But that's enough, I think, to visualize the whole plane because if we take a fraction of those, we fill in that inside. But also, we take 3 times them and 4-- any number times these particular vectors, and they'll go outside the part I've drawn.

So I've drawn a part of the row space of the matrix, combinations of the rows. And I've marked out 20 combinations-- 1, 2, 3, 4-- maybe only 16. Yeah, but it's enough for us to imagine what the whole, filled-in plane would look like. OK, so that's the other key space, the row space. Now, only two more spaces to go in this example.

The null space, what does that mean? Again, it's a space of vectors. And in this, the null space of the matrix has all the vectors x that give Ax equals 0. If you know all those, then that's a natural bunch of vectors. So here, we have to solve Ax equals 0. So you see the A matrix, the 1, 2, 3, 4, 5, 6. And you see the x vector. And you see the multiple. So I did multiplication. And I've got the dot products there. And now I'm setting it to 0. Do you see any x1 and x2 and x3 that work?

I have two equations, both with 0 on the right-hand side. And they have a solution because I have three unknowns, two equations. So there should be some freedom here. And there, at the bottom line of this slide, it gives the answer. If I take 1 of the first column minus 2 of the second plus 1 of the third, I get 0. And that means I found something in the null space. Do you see that, going back to the matrix? If I add 1,4 and 3, 6, that gives me 410. And that's just twice 2, 5. So that line of vectors, through 1 minus 2, 1, is the null space. If you've got the null space-- and that's what elimination in linear algebra, finds a solution.

One more space to go. And it is a little special because it only contains the 0 vector. So again, this is the null space of A transpose. So I flipped those rows and made them columns again, in A transpose. Now I have three equations and only two unknowns. And there could be solutions if I had three equations and two unknowns. There could. There's always the 0 solution.

And in this case, that's all there is. That's all there is. Those two columns are in-- 1 2, 3, that column is a vector in 3D. 4, 5, 6 points another way. And any combination is just going to have some part of 1, 2, 3 and some part of 4, 5, 6. The only way to get to the 0 vector is to have 0 of each. So the only solution to these three equations with two unknowns is 0 and 0. So the null space of A transpose in this example was pretty small. It's got one vector in it though, the 0 vector. But that's all.

So now we've done the four spaces. May I bring back the picture? And you'll see how they fit again. So this is the column space. Let's review. The column space had two columns in it. It didn't need all three columns of A. It just had two columns and their combinations, which picked up the third column. So it was dimension two. This space, this matrix, has-- the letter r stands for the rank of the matrix. And that's the number of independent columns. And in this case, you remember there were two. The third one wasn't important. And it also-- this is a beautiful fact-- that same r is the number of independent rows. And you remember this matrix had just two rows and they were independent. So the rank is 2.

We can find the rank from the column space, or we can find the rank from the row space And a beautiful bit of mathematics is that the number of vectors you need is the same. And then the null space of A, I think we found one vector perpendicular to the rows. And the null space of A transpose might have been 0 in that case, or vice versa. Anyway, we have found all the numbers that go with the example. But here, you're seeing not an example, but the big picture. This is the fundamental picture of linear algebra. And I've learned that in reading that section of my textbook, students really bring things together.

OK, so let's see. What do I want to say next about these four subspaces? Apart from recommending that you take some matrix yourself and figure out the four subspaces? So here's the fact that I've been talking about, the column rank equaling the row rank. And that's a beautiful but not obvious fact. If you had a 50 by 70 matrix, well, just finding how many independent columns and how many independent rows would be a challenge. But the numbers would be the same. The number of independent columns equals the number of independent rows.

That's just linear algebra magic. And that number, we call r. So that measures the real size of the matrix And every textbook has to prove that. And I like to build it out of this factorization. For me, factorizations are fundamental to linear algebra, for everybody, really. I'll come back in on another day, maybe to talk about that factorization.

But the rank is 2, here. Here is our friendly matrix, A. And we divided it into-- we took the independent columns, 1, 4 and 2, 5 as the C matrix. And then the R matrix, we just filled in to get the right answer. And 2 was the number of independent columns. And 2 is the number of independent rows. So this A equals CR gives a beautiful proof of that dimension theorem, dimension of row space equal dimension of column space. So that's just a nice part of the theory.

OK, now we are going to use the four subspaces. Everybody now knows, the row space and the column space have the same dimension, r. And if I take something in the row space, multiply by A, it's in the column space, of course. Everything, when I multiply by A, fills the column space. And for that part of the problem, from row space to column space, from space of dimension r to a space of dimension r, the matrix is invertible.

And I give a proof that you don't need. But that's the sign of an invertible matrix, when All vectors go to-- all different vectors here go to different vectors here. And that's from row space to column space. That's a key-- that's the heart of the matrix, is the part that actually does something non-0, that takes the row space to the column space.

And then you can think about the matrix that brings it back, just undoes this. And that's called the pseudoinverse. And its symbol is A with a plus sign or a dagger. So that's a newish, less familiar idea, that because the matrix A, between row space and column space-- the null space is not in the picture for this slide-- just between row space and column space, that's a perfect match, those two spaces. So there's a matrix that brings me backwards, an inverse matrix. But it's called the pseudo inverse because I'm ignoring any null spaces. It's not a true inverse because a true inverse shouldn't have a null space there.

OK, now, application. So I've discussed the four subspaces. And I suggest you figure out what they might be for some other matrix. But now, I want to use that idea in the most important type of application for statistics. It even has two names. Statistics people call it regression. And linear algebra people call it least squares. And that's the problem there. Find the vector x that makes Ax minus b as short as possible minimizes the length. If there's a solution to Ax equals b, then we take that x and we get Ax minus b would be 0. But when we have null spaces and things around, we may not be able to get Ax equal b.

But what we can get, and do get, is that equation A transpose times A times x equals A transpose times b. In other words, you could see actually, by calculus if you took the derivative of Ax minus b squared because you always take derivatives if you're minimizing something, and set them to 0, it would lead to that equation. That square brings in an A transpose. And that's the equation you get.

So that's the equation that least squares problems have to solve. And it's one of the fundamental equations of linear algebra. And so what I'm saying down here is, there's one x hat. Now, I'm calling this solution here x hat because I'm not solving Ax equals b. I'm solving this easier. I made more solutions or created some solutions when I multiply both sides by A transpose. So the solutions now, are called x hat.

And there's one solution in the null space. There is one solution, x hat, in the row space because the row space and the columns-- we have-- this is r equations and r unknowns. And it's got a unique solution, x hat. That x hat is the goal. That's the answer to the regression problem, the least squares problem. X hat makes Ax minus b as small as we can make it. That's the least-squares problem, least squares, smallest square. And so that's the one we're after.

O, just to finish up on the pseudoinverse, which is a cool idea, let me come back to the main slide of the whole talk, the four subspaces, and ask what the pseudoinverse does. You remember what the matrix A does. It takes a vector in the row space, multiply by A, you get something in the column space. Something in the null space, you multiply by A and you get 0.

Now we want to go backwards. So the beauty is that this part, which is what we want to just reverse, those have the same dimension. So we're fundamentally looking at a square matrix which has an inverse when we're up top. So that's what the pseudoinverse is. It takes the y back to x. And what does it do to these Sends them to 0 just the way A transpose does. So the pseudoinverse gives the correct inverse up here where the real action is. And down here, where vectors are sent to 0-- can't bring them back from 0. They're dead. So every vector here goes to 0.

So that's the pseudoinverse. The four subspaces are just reversed, right to left. And it's a useful idea-- not the most basic idea. It doesn't show up in all courses. But when we had the four subspaces here, it was so easy to tell you what the pseudoinverse is. So it just takes-- it's so neat with this figure. It takes each vector in the column space back to the vector in the row space where it came from. So that's A plus the pseudoinverse.

OK, thank you. That's the four subspaces of any matrix. And you'll remember that if a matrix is invertible, square and invertible, these two spaces are gone, and we simply have-- the row space and the column space are both the full space. And we have a perfect one-to-one exchange. And luckily, a lot of matrices are invertible. But this covers the case of matrices. Invertible or not, square or not, every matrix has got those four fundamental subspaces. Thank you.