18.06 Linear Algebra, Spring 2005
Transcript – Lecture 8

OK, when the camera says, we'll start.

You want to give me a signal? OK, this is lecture eight in linear algebra, and this is the lecture where we completely solve linear equations.

So Ax=b. That's our goal. If it has a solution. It certainly can happen that there is no solution. We have to identify that possibility by elimination. And then if there is a solution we want to find out is there only one solution or are -- is there a whole family of solutions, and then find them all. OK. Can I use as an example the same matrix that I had last time when we were looking for the null space.

So the, the matrix has rows 1 2 2 2, 2 4 6 8, and the third row -- you remember the main point was the third row, 3 6 8 10, is the sum of row one plus row two. In other words, if I add those left-hand sides, I get the third left-hand side. So you can tell me right away what elimination is going to discover about the right-hand sides. What's -- there is a condition on b1, b2, and b3 for this system to have a solution. Most cases -- if I took these numbers to be one -- 5, and 17, there would not be a solution. In fact, if I took those first numbers to be 1 and 5, what is the only b3 that would be OK? Six.

If the left-hand -- if these left-hand sides add up to that, then B -- I need b1 plus b2 to equal b3.

Let's just see how elimination discovers that.

But we can see it coming, right? That if -- let me say it in other words.

If some combination on the left-hand side gives all 0s then the same combination on the right-hand side must give 0.

OK. So let me take that example and write down instead of copying out all the plus signs, let me write down the matrix. 1 2 2 2, 2 4 6 8, and that 6 3 8 10, where the third row is the sum of the first two rows. Now how do we deal with the right-hand side? That's -- we want to do the same thing to the right-hand side that we're doing to these rows on the left side, so we just tack on the right-hand side as another vector, another column.

So this is the augmented matrix.

It's, it's the matrix A with the vector b tacked on.

In Matlab, that's all you would need to type. OK. So we do elimination on that.

Can we just do elimination quickly? The first pivot is fine, I subtract two of this away from this, three of this away from this, so I have 1 2 2 2 b1.

Two of those away will give me 0 0 2 and 4, and that was b2 minus two b1. I, I have to do the same thing to that third, that last column.

And then three of these away from this gave me 0 0 2 4 b3 minus three b1s. So that's the, that's elimination with the first column completed.

We move on. There's the first pivot still.

Here is the second pivot. We're always remembering, now, these are then going to be the pivot columns.

And let me get the final result -- well, let me -- can I do it by eraser? We're capable of subtracting this row from this row, just by -- that'll knock this out completely and give me the row of 0s, and on the right-hand side, when I subtract this away from this, what do I have? I think I have b3 minus a b2, and I had minus three b1s.

This is going to, it's going to be a minus a b1.

Oh yeah that's exactly what I expect.

So now the -- what's the last equation? The last equation, this represented by this zero row, that last equation is, says 0 equals b3 minus b2 minus b1. So that's the condition for solvability. That's the condition on the right-hand side that we expected.

It says that b1+b2 has to match b3, and if our numbers happen to have been 1, 5, and 6 -- so let me take, suppose b is 1 5 6. That's an OK b.

And when I do this elimination, what will I have? The b1 will still be a 1. b2 would be 5 minus 2, this would be a 3. 5 -- my 6 minus 5 minus 1, this will be -- this is the main point -- this will be a 0, thanks. OK.

So the last equation is OK now. And I can proceed to solve the two equations that are really there with four unknowns.

OK, I, I, I want to do that, so this, this b is OK.

It allows a solution. We're going to be, naturally, interested to keep track what are the conditions on b that make the equation solvable. So let me write down what we already see before I continue to solve it. Let me first -- solvability, solvability. So which -- so this is the condition on the right-hand sides.

And what is that condition? This is solvability always of Ax=b. So Ax=b is solvable -- well, actually, we had an answer in the language of the column space. Can you remind me what that answer is? That, that was like our answer from earlier lecture. b had to be in the column space. Solvable if -- when -- exactly when b is in the column space of A.

Right? That just says that b has to be a combination of the columns, and of course that's exactly what the equation is looking for.

So that -- now I want to answer it -- the same answer but in different language. Another way to answer this -- if a combination of the rows of A gives the zero row, and this was an example where it happened, some combination of the rows of A produced the zero row -- then what's the requirement on b? Since we're going to do the same thing to both sides of all equations -- the same combination of the components of b has to give 0. Right? That's -- so if there's a combination of the rows that gives the zero row, then the same combination of the entries of b must give 0.

And this isn't the zero row, that's the zero number. OK. Tthis is another way of saying -- and it is not immediate, right, that these two statements are equivalent. But somehow they must be, because they're both equivalent to the solvability of the system. OK.

So we've got this, this sort of -- like question zero is, does the system have a solution? OK, I'll come back to discuss that further.

Let's go forward when it does. When there is a solution.

And so what's our job now? Abstractly we sit back and we say, OK, there's a solution, finished.

It exists. But we want to construct it.

So what's the algorithm, the sequence of steps to find the solution? That's what I -- and of course the quiz and the final, I'm going to give you a system Ax=b and I'm going to ask you for the solution, if there is one. And so this algorithm that you want to follow. OK, let's see.

So what's the -- so now to find the complete solution to Ax=b.

OK. Let me start by finding one solution, one particular solution.

I'm expecting that I can, because my system of equations now, that last equation is zero equals zero, so that's all fine.

I really have two equations -- actually I've got four unknowns, so I'm expecting to find not only a solution but a whole bunch of them. But let's just find one.

So step one, a particular solution, x particular. How do I find one particular solution? Well, let me tell you how I, how I find it. So this is -- since there are lots of solutions, you could have your own way to find a particular one. But this is a pretty natural way. Set all free variables to zero.

Since those free variables are the guys that can be anything, the most convenient choice is zero.

And then solve Ax=b for the pivot variables.

So what does that mean in this example? Which are the free variables? Which, which are the variables that we can assign freely and then there's one and only one way to find the pivot variables? They're x2 and -- so x2 is zero, because that's in a column

without a pivot, the second column has no pivot. And the -- what's the other one? The fourth, x4 is zero.

Because that, those are the, the free ones. Those are in the columns with no pivots. So you see what my -- so when I knock -- when x2 and x4 are zero, I'm left with the -- what I left with here? I'm just left with -- see, now I'm not using the two free columns.

I'm only using the pivot columns.

So I'm really left with x1 -- the first equation is just x1 and two x3s should be the right-hand side, which we picked to be a one. And the second equation is two x3s, as it happened, turned out to be, three. I just write it again here with the x2 and the x4 knocked out, since we're set them to zero. And you see that we're back in the normal case of having back -- where back substitution will do it.

So x3 is three halves, and then we go back up and x1 is one minus two x3. That's probably minus two.

Good. So now we have the solution, x particular is the vector minus two zero three halves zero. OK, good.

That's one particular solution, and we should and could plug it into the original system. Really if -- on the quiz, please, it's a good thing to do. So we did all this, these, row operations, but this is supposed to solve the original system, and I think it does. OK.

So that's x particular which we've got.

So that's like what's new today.

The particular solution comes -- first you check that you have zero equals zero, so you're OK on the last equations. And then you set the free variables to zero, solve for the pivot variables, and you've got a particular solution, the particular solution that has zero free variables.

OK. Now -- but that's only one solution, and now I'm looking for all.

So how do I find the rest? The point is I can add on x -- anything out of the null space. We know how to find the vectors in the null space -- because we did it last time, but I'll remind you what we got.

And then I'll add. So the final result will be that the complete solution -- this is now the complete guy -- the complete solution is this one particular solution plus any, any vector, all different vectors out of the null space. xn, OK.

Well why, why this pattern, because this pattern shows up through all of mathematics, because it shows up everywhere we have linear equations. Let me just put here the, the reason. A xp, so that's x particular, so what does Ax particular give? That gives the correct right-hand side b. And what does A times an x in the null space give? Zero. So I add, and I put in parentheses. So xp plus xn is b plus zero, which is b. So -- oh, what I saying? Let me just say it in words. If I have one

solution, I can add on anything in the null space, because anything in the null space has a zero right-hand side, and I still have the correct right-hand side B.

So that's my system. That's my complete solution.

Now let me write out what that will be for this example.

So in this example, x general, x complete, the complete solution, is x particular, which is minus two zero three halves zero, with those zeroes in the free variable, plus -- you remember there were the special solutions in the null space that had a one in the free variables -- or one and zero in the free variables, and then we filled in to find the others? I've forgotten what they were, but maybe it was that. That was a special solution, and then there was another special solution that had that free variable zero and this free variable equal one, and I have to fill those in. Let's see, can I remember how those fill in? Maybe this was a minus two and this was a two, possibly? I think probably that's right. I'm not -- yeah.

Does that look write to you? I would have to remember what are my equations. Can I, rather than go way back to that board, let me remember the first equation was two x3 plus two x4 equaling zero now, because I'm looking for the guys in the null space.

So I set x4 to be one and the second equation, that I didn't copy again, gave me minus two for this and then -- yeah, so I think that's right.

Two minus four and two gives zero, check.

OK. Those were the special solutions. What do we do to get the complete solution? How do I get the complete solution now? I multiply this by anything, c1, say, and I multiply this by anything -- I take any combination. Remember that's how we described the null space? The null space consists of all combinations of -- so this is xn -- all combinations of the special solutions.

There were two special solutions because there were two free variables. And we want to make that count -- carefully now. Just while I'm up here.

So there's, that's what the -- that's the kind of answer I'm looking for. Is there a constant multiplying this guy? Is there a free constant that multiplies x particular? No way.

Right? x particular solves A xp=b.

I'm not allowed to multiply that by three.

But Axn, I'm allowed to multiply xn by three, or add to another xn, because I keep getting zero on the right. OK.

So, so again, xp is one particular guy.

xn is a whole subspace. Right? It's one guy plus, plus anything from a subspace.

Let me draw it. Let me try to -- oh. I want to draw, I want to graph all this -- I want to, I want to plot all solutions.

Now x. So what dimension I in? This is a unfortunate point. How many components does x have? Four.

There are four unknowns. So I have to draw a four dimensional picture on this MIT cheap blackboard.

OK. So here we go.

x1 -- Einstein could do it, but, this, this is -- those are four perpendicular axes in -- representing four dimensional space. OK.

Where are my solutions? Do my solutions form a subspace? Does the set of solutions to Ax=b form a subspace? No way.

What does it actually look like, though? A subspace is in this picture. This part is a subspace, right? That part is some, like, two dimensional, because I've got two parameters, so it's -- I'm thinking of this null space as a two dimensional subspace inside R^4. Now I have to tell you and will tell you next time, what does it mean to say a subspace, what's the dimension of a subspace. But you see what it's going to be. It's the number of free independent constants that we can choose.

So somehow there'll be a two dimensional subspace, not a line, and not a three dimensional plane, but only a two dimensional guy. But it's doesn't go through the origin because it goes through this point.

So there's x particular. x particular is somewhere here.

x particular. So it's somehow a subspace -- can I try to draw it that way? It's a two dimensional subspace that goes through x particular and then onwards by -- so there's x particular, and I added on xn, and there's x. There's x=xp+xn. But the xn was anywhere in this subspace, so that filled out a plane. It's a subspace -- it's not a subspace, what I saying? It's like a flat thing, it's like a subspace, but it's been shifted, away from the origin. It doesn't contain zero. OK. Thanks.

That's the picture, and that's the algorithm.

So the algorithm is just go through elimination and, find the particular solution, and then find those special solutions. You can do that. Let me take our time here in the lecture to think, about the bigger picture. So let me think about -- so this is my pattern. Now I want to think -- I want to ask you about a question -- I want to ask you some questions.

So when I mean think bigger, I mean I'll think about an m by n matrix A of rank r. OK.

What's our definition of rank? Our current definition of rank is number of pivots. OK.

First of all, how are these numbers related? Can you tell me a relation between r and m? If I have m rows in the matrix and R pivots, -- then I certainly know, always -- what relation do I know between r and m? r is less or equal, right? Because I've got m rows, I can't have more than m pivots, I might have m and I might have fewer.

Also, I've got n columns. So what's the relation between r and n? It's the same, less or equal, because a column can't have more than one pivot. So I can't have more than n pivots altogether. OK, OK.

So I have an m by n matrix of rank r.

And I always know r less than or equal to m, r less than or equal to n. Now I'm specially interested in the case of full rank, when the rank r is as big as it can be. Well, I guess I've got two separate possibilities here, depending on what these numbers m and n are. So let me talk about the case of full column rank. And by that I mean r=n.

And I want to ask you, what does that imply about our solutions? What does that tell us about the null space? What does that tell us about, the complete solution? OK, so what does that mean? So I want to ask you, well, OK, if the rank is n, what does that mean? That means there's a pivot in every column. So how many pivot variables are there? n.

All the columns have pivots in this case. So how many free variables are there? None at all. So no free variables. r=n, no free variables. So what does that tell us about what's going to happen then in our, in our little algorithms? What will be in the null space? The null space of A has got what in it? Only the zero vector.

There are no free variables to give other values to.

So the null space is only the zero vector.

And what about our solution to Ax=b? Solution to Ax=b? What, what's the story on that one? So now that's coming from today's lecture. The solution x is -- what's the complete solution? It's just x particular, right? If, if, if there is an x, if there is a solution. It's x equal x particular.

There's nothing -- you know, there's just one solution.

If there's one at all. So it's unique solution -- unique means only one -- unique solution if it exists, if it exists. In other words, I would say -- let me put it a different way.

There're either zero or one solutions.

This is all in this case r=n. So I'm -- because many, many applications in reality, the columns will be what I'll later call independent. And we'll have, nothing to look for in the null space, and we'll only have particular solutions. OK.

Everybody see that possibility? But I need an example, right? So let me create an example.

What sort of a matrix -- what's the shape of a matrix that has full column rank? So can I squeeze in an, an example here? If it exists.

Let me put in an example, and it's just the right space to put in an example. Because the example will be like tall and thin. It will have -- well, I mean, here's an example, one two six five, three one one one. Brilliant example.

OK. So there's a matrix A, and what's its rank? What's the rank of that matrix? How many pivots will I find if I do elimination? Two, right? Two. I see a pivot there -- oh certainly those two columns are headed off in different directions.

When I do elimination, I'll certainly get another pivot here, fine, and I can use those to clean out below and above. So the -- actually, tell me what its row reduced row echelon form would be.

Can you carry that, that elimination process to the bitter end? So what do, what does that mean? I subtract a multiple of this row from these rows. So I clean up, all zeros there. Then I've got some pivot here.

What do I do with that? I go subtract it below and above, and then I divide through, and what's R for that example? Maybe I can -- you'll allow me to put that just here in the next board.

What's the row reduced echelon form, just out of practice, for that matrix? It's got ones in the pivots.

It's got the identity matrix, a little two by two identity matrix, and below it all zeros. That's a matrix that really has two independent rows, and they're the first two, actually. The first two rows are independent. They're not in the same direction. But the other rows are combinations of the first two, so -- is there always a solution to Ax=b? Tell me what's the picture here? For this matrix A, this is a case of full column rank.

The two columns are -- give two pivots.

There's nothing in the null space.

There's no combination of those columns that gives the zero column except the zero zero combination.

So there's nothing in the null space.

But is there always a solution to A X equal B? What's up with A X equal B? I've got four, four equations here, and only two Xs.

So the answer is certainly no. There's not always a solution.

I may have zero solutions, and if I make a random choice, I'll have zero solutions. Or if I make a great particular choice of the right-hand side, which just happens to be a combination of those two guys -- like tell me one right-hand side that would have a solution. Tell me a right-hand side that would have a solution. Well, 0 0 0 0, OK. No prize for that one.

Tell me another one. Another right-hand side that has a solution would be 4 3 7 6. I could add the two columns. Right? What would be the total complete solution if the right-hand side was 4 3 7 6? There would be the particular solution one one, one of that column plus one of that, and that would be the only solution. So there would be -- x particular would be one one in the case when the right side is the sum of those two columns, and that's it.

So that would be a case with one solution.

OK. That, this is the typical setup with full column rank. Now I go to full row rank.

You see the sort of natural symmetry of this discussion.

Full row rank means r=m. So this is what I'm interested in now, r=m. OK, what's up with that? How many pivots? m.

So what happens when we do elimination in that case? I'm going to get m pivots. So every row has a pivot, right? Every row has a pivot.

Then what about solvability? What about this business of -- for which right-hand sides can I solve it? So that's my question. I can solve Ax=b for which right-hand sides? Do you see what's coming? I do elimination, I don't get any zero rows.

So there aren't any requirements on b.

I can solve Ax=b for every b. I can solve Ax=b for every right-hand side. So this is the existence, exists a solution. Now tell me, so the, u- u- so every row has a pivot in it.

So how many free variables are there? How many free variables in this case? If I had n variables to start with, how many are used up by pivot variables? r, which is m. So I'm left with, left with n-r free variables.

OK. So this case of full row rank I can always solve, and then this tells me how many variables are free, and this is of course n-m. This is n-m free variables. Can I do an example? You know, the best way for me to do an example is just to transpose that example. So let me take, let me take that matrix that had column one two six five and make it a row. And let me take three one one one as the second row. And let me ask you, this is my matrix A, what's its rank? What's the rank of that matrix? Sorry to ask, but not sorry really, because we're just getting the idea of rank. What's the rank of that matrix? Two, exactly, two.

There will be two pivots. What will the row reduced echelon form be? Anybody know that one? Actually, tell me not only -- you have to tell me not only the, there'll be two pivots but which will be the pivot columns.

Which columns of this matrix will be pivot columns? So the first column is fine, and then I go on to the next column, and what do I get? Do I get a second pivot out of -- will I get a second pivot in this position? Yes.

So the pivots, when I get all the way to R, will be there. And here will be some numbers.

This is the part that I previously called F. This is the part that -- the pivot columns in R will be the identity matrix. There are no zero rows, no zero rows, because the rank is two.

But there'll be stuff over here.

And that will, enter the special solutions and the null space. OK.

So this is a typical matrix with r=m smaller than n.

Now finally I've got a space here for r=m=n.

I'm off in the corner here with the most important case of all.

So what's up with this matrix? So let me give an example.

OK, brilliant example, 1 2 3 1.

Tell me what -- how do I describe a matrix that has rank r=m=n? So the matrix is square, right, it's a square matrix. And if I know its rank is -- it's full rank, now.

I don't have to say full column rank or full row rank -- I just say full rank, because the count, column count and the row count are the same, and the rank is as big as it can be.

And what kind of a matrix have I got? It's invertible. So that's exactly the invertible matrices. r=m=n means the -- what's the row echelon form, the reduced row echelon form, for an invertible matrix? For a square, nice, square, invertible matrix? It's I. Right.

So you see that the, the good matrices are the ones that kind of come out trivially in R. You reduce them all the way to the identity matrix. What's the null space for this, for this matrix? Can I just hammer away with questions? What's the null space for this matrix? The null space of that matrix is the zero vector only.

The zero vector only. What are the conditions to solve Ax=b? Which right-hand sides b are OK? If I want to solve Ax=b for this example, so A is this, b is b1 b2, what are the conditions on b1 and b2? None at all, right.

So this is the case, this is the case where I can solve -- so I've coming back here, I can -- since the rank equals m, I can solve for every b.

And since the rank is also n, there's a unique solution.

Let me summarize the whole picture here.

Here's the whole picture. I could have r=m=n.

This is the case where this is the identity matrix. And this is the case where there is one solution.

That's the square invertible chapter two case. Now we're into chapter three. We could have r=m smaller than n. Now that's what we had over there, and the row echelon form looked like the identity with some zero rows. And that was the case where there are zero or one solution. When I say solution I mean to Ax=b. So this case, there's always one. This case there's zero or one.

And now let me take the case of full column rank, but some, extra rows. So now R has -- well, the identity -- I'm almost tempted to write the identity matrix and then

F, but that isn't necessarily right. I have -- is that right? Am I getting this correct here? Oh, I'm not! My God! This is the case R equals n, the columns, the columns are, are OK. That's the case that was on that board, r=n, full column rank.

Now I want the case where m is smaller than n and I've got extra columns. OK.

There we go. So this is now the case of full row rank, and it looks like I F except that I can't be sure that the pivot columns are the first columns.

So the I and the F, could be partly mixed into the I. Can I write that with just like that? So the F could be sort of partly into the I if the first columns weren't the pivot columns. Now how many solutions in this case? There's always a solution.

This is the existence case. There's always a solution. We're not getting any zero rows.

There are no zero rows here. So there's always either one or infinitely many solutions. OK.

Actually, I guess there's always an infinite number, because we always have some null space to deal with.

Then the final case is where r is smaller than m and smaller than n. OK.

Now that's the case where R is the identity with some free stuff but with some zero rows too.

And that's the case where there's either no solution -- because we didn't get a zero equals zero for some bs, or infinitely many solutions. OK.

Do you -- this board really summarizes the lecture, and this sentence summarizes the lecture.

The rank tells you everything about the number of solutions.

That number, the rank r, tells you all the information except the exact entries in the solutions.

For that you go to the matrix. OK, good.

Have a great weekend, and I'll see you on Monday.