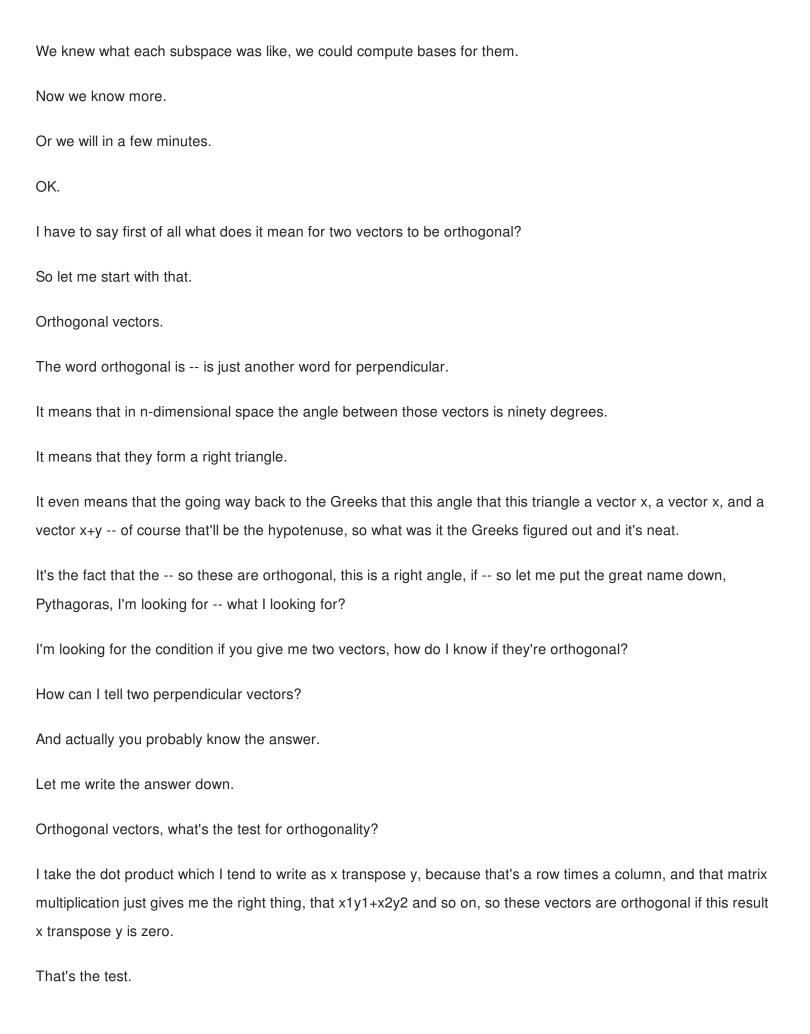OK. cameras are rolling. This is lecture fourteen, starting a new chapter.

Chapter about orthogonality.

What it means for vectors to be orthogonal.

What it means for subspaces to be orthogonal.

What it means for bases to be orthogonal.

So ninety degrees, this is a ninety-degree chapter.

So what does it mean -- let me jump to subspaces.

Because I've drawn here the big picture.

This is the 18.06 picture here.

And hold it down, guys.

So this is the picture and we know a lot about that picture already.

We know the dimension of every subspace.

We know that these dimensions are r and n-r. We know that these dimensions are r and m-r. What I want to show now is what this figure is saying, that the angle -- the figure is just my attempt to draw what I'm now going to say, that the angle between these subspaces is ninety degrees.

And the angle between these subspaces is ninety degrees.

Now I have to say what does that mean?

What does it mean for subspaces to be orthogonal?

But I hope you appreciate the beauty of this picture, that that those subspaces are going to come out to be orthogonal.

Those two and also those two.

So that's like one point, one important point to step forward in understanding those subspaces.

We knew what each subspace was like, we could compute bases for them.

Now we know more.

Or we will in a few minutes.

OK.

I have to say first of all what does it mean for two vectors to be orthogonal?

So let me start with that.

Orthogonal vectors.

The word orthogonal is -- is just another word for perpendicular.

It means that in n-dimensional space the angle between those vectors is ninety degrees.

It means that they form a right triangle.

It even means that the going way back to the Greeks that this angle that this triangle a vector x, a vector x, and a vector x+y -- of course that'll be the hypotenuse, so what was it the Greeks figured out and it's neat.

It's the fact that the -- so these are orthogonal, this is a right angle, if -- so let me put the great name down, Pythagoras, I'm looking for -- what I looking for?

I'm looking for the condition if you give me two vectors, how do I know if they're orthogonal?

How can I tell two perpendicular vectors?

And actually you probably know the answer.

Let me write the answer down.

Orthogonal vectors, what's the test for orthogonality?

I take the dot product which I tend to write as x transpose y, because that's a row times a column, and that matrix multiplication just gives me the right thing, that x1y1+x2y2 and so on, so these vectors are orthogonal if this result x transpose y is zero.

That's the test.

OK.

Can I connect that to other things?

I mean -- it's just beautiful that here we have we're in n dimensions, we've got a couple of vectors, we want to know the angle between them, and the right thing to look at is the simplest thing that you could imagine, the dot product.

OK. Now why?

So I'm answering the question now why -- let's add some justification to this fact, that's the test.

OK, so Pythagoras would say we've got a right triangle, if that length squared plus that length squared equals that length squared.

OK, can I write it as x squared plus y squared equals x plus y squared?

Now don't, please don't think that this is always true.

This is only going to be true in this -- it's going to be equivalent to orthogonality.

For other triangles of course, it's not true.

For other triangles it's not.

But for a right triangle somehow that fact should connect to that fact.

Can we just make that connection?

What's the connection between this test for orthogonality and this statement of orthogonality?

Well, I guess I have to say what is the length squared?

So let's continue on the board underneath with that equation.

Give me another way to express the length squared of a vector.

And let me just give you a vector.

The vector one, two, three.

That's in three dimensions.

What is the length squared of the vector x equal one, two, three?

So how do you find the length squared?

Well, really you just, you want the length of that vector that goes along one -- up two, and out three -- and we'll come back to that right triangle stuff.

The length squared is exactly x transpose x.

Whenever I see x transpose x, I know I've got a number that's positive.

It's a length squared unless it -- unless x happens to be the zero vector, that's the one case where the length is zero.

So right -- this is just x1 squared plus x2 squared plus so on, plus xn squared.

So one -- in the example I gave you what was the length squared of that vector one, two, three?

So you square those, you get one, four and nine, you add, you get fourteen.

So the vector one, two, three has length fourteen.

So let me just put down a vector here.

Let x be the vector one, two, three.

Let me cook up a vector that's orthogonal to it.

So what's the vector that's orthogonal?

So right down here, x squared is one plus four plus nine, fourteen, let me cook up a vector that's orthogonal to it, we'll get right that that -- those two vectors are orthogonal, the length of y squared is five, and x plus y is one and two making three, two and minus one making one, three and zero making three, and the length of this squared is nine plus one plus nine, nineteen.

And sure enough, I haven't proved anything.

I've just like checked to see that my test x transpose y equals zero, which is true, right?

Everybody sees that x transpose y is zero here?

That's maybe the main point.

That you should get really quick at doing x transpose y, so it's just this plus this plus this and that's zero.

And sure enough, that clicks with fourteen plus five agreeing with nineteen.

Now let me just do that in letters.

So that's y transpose y.

And this is x plus y transpose x plus y.

OK. So I'm looking, again, this isn't always true.

I repeat.

This is going to be true when we have a right angle.

And let's just -- well, of course, I'm just going to simplify this stuff here.

There's an x transpose x there.

And there's a y transpose y there.

And there's an x transpose y.

And there's a y transpose x.

I knew I could do that simplification because I'm just doing matrix multiplication and I've just followed the rules.

OK.

So x transpose x-s cancel.

Y transpose y-s cancel.

And what about these guys?

What can you tell me about the inner product of x with y and the inner product of y with x?

Is there a difference?

I think if we -- while we're doing real vectors, which is all we're doing now, there isn't a difference, there's no difference.
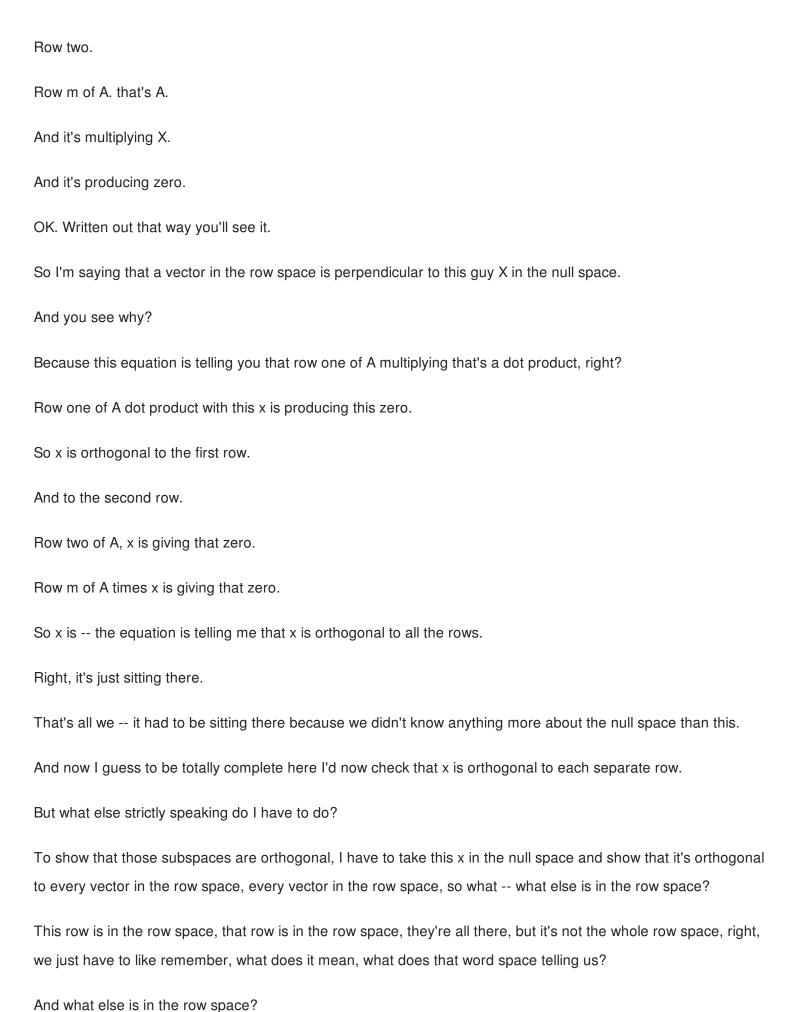
If I take x transpose y, that'll give me zero, if I took y transpose x I would have the same x_1y_1 and x_2y_2 and x_3y_3, it would be the same, so this is -- this is the same as that, this is really I'll knock that guy out and say two of these.

So actually that's the -- this equation boiled down to this thing being zero.

Right? Everything else canceled and this equation boiled down to that one.

So that's really all I wanted.

I just wanted to check that Pythagoras for a right triangle led me to this -- of course I cancel the two now.

No problem.

To x transpose y equals zero as the test.

Fair enough.

OK.

You knew it was coming.

The dot product of orthogonal vectors is zero.

It's just -- I just want to say that's really neat.

That it comes out so well.

All right.

Now what about -- so now I know if two -- when two vectors are orthogonal.

By the way, what about if one of these guys is the zero vector?

Suppose x is the zero vector, and y is whatever.

Are they orthogonal?

Sure.

In math the one thing about math is you're supposed to follow the rules.

So you're supposed to -- if x is the zero vector, you're supposed to take the zero vector dotted with y and of

course you always get zero.

So just so we're all sure, the zero vector is orthogonal to everybody.

But what I want to -- what I now want to think about is subspaces.

What does it mean for me to say that some subspace is orthogonal to some other subspace?

So OK.

Now I've got to write this down.

So because we're defining definition of subspace S is orthogonal so to subspace let's say T, so I've got a couple of subspaces.

And what should it mean for those guys to be orthogonal?

It's just sort of what's the natural extension from orthogonal vectors to orthogonal subspaces?

Well, and in particular, let's think of some orthogonal subspaces, like this wall.

Let's say in three dimensions.

So the blackboard extended to infinity, right, is a -- is a subspace, a plane, a two-dimensional subspace.

It's a little bumpy but anyway, it's a -- think of it as a subspace, let me take the floor as another subspace.

Again, it's not a great subspace, MIT only built it like so-so, but I'll put the origin right here.

So the origin of the world is right there.

OK.

Thereby giving linear algebra its proper importance in this.

OK. So there is one subspace, there's another one.

The floor.

And are they orthogonal?

What does it mean for two subspaces to be orthogonal and in that special case are they orthogonal?

All right.

Let's finish this sentence.

What does it mean means we have to know what we're talking about

here. So what would be a reasonable idea of orthogonal?

Well, let me put the right thing up.

It means that every vector in S, every vector in S, is orthogonal to -- what I going to say?

Every vector in T.

That's a reasonable and it's a good and it's the right definition for two subspaces to be orthogonal.

But I just want you to see, hey, what does that mean?

So answer the question about the -- the blackboard and the floor.

Are those two subspaces, they're two-dimensional, right, and we're in R^3. It's like a xz plane or something and a xy plane.

Are they orthogonal?

Is every vector in the blackboard orthogonal to every vector in the floor, starting from the origin right there?

Yes or no?

I could take a vote.

Well we get some yeses and some noes.

No is the answer.

They're not.

You can tell me a vector in the blackboard and a vector in the floor that are not orthogonal.

Well you can tell me quite a few, I guess.

Maybe like I could take some forty-five-degree guy in the blackboard, and something in the floor, they're not at ninety degrees, right?

In fact, even more, you could tell me a vector that's in both the blackboard plane and the floor plane, so it's certainly not orthogonal to itself.

So for sure, those two planes aren't orthogonal.

What would that be?

So what's a vector that's -- in both of those planes?

It's this guy running along the crack there, in the intersection, the intersection.

A vector, you know -- if two subspaces meet at some vector, well then for sure they're not orthogonal, because that vector is in one and it's in the other, and it's not orthogonal to itself unless it's zero.

So the only I mean so orthogonal is for me to say these two subspaces are orthogonal first of all I'm certainly saying that they don't intersect in any nonzero vector.

But also I mean more than that just not intersecting isn't good

enough. So give me an example, oh, let's say in the plane, oh well, when do we have orthogonal subspaces in the plane?

Yeah, tell me in the plane, so we don't -- there aren't that many different subspaces in the plane.

What what have we got in the plane as possible subspaces?

The zero vector, real small.

A line through the origin.

Or the whole plane.

OK. Now so when is a line through the origin orthogonal to the whole plane?

Never, right, never.

When is a line through the origin orthogonal to the zero subspace?

Always.

Right. When is a line through the origin orthogonal to a different line through the origin?

Well, that's the case that we all have a clear picture of, they -- the two lines have to meet at ninety degrees.

They have only the -- so that's like this simple case I'm talking about.

There's one subspace, there's the other subspace.

They only meet at zero.

And they're orthogonal.

OK. Now. So we now know what it means for two subspaces to be orthogonal.

And now I want to say that this is true for the row space and the null space.

OK. So that's the neat fact.

So row space is orthogonal to the null space.

Now how did I come up with that?

But you see the rank it's great, that means that these -- that these subspaces are just the right things, they're just cutting the whole space up into two perpendicular subspaces.

OK.

So why? Well, what have I got to work with?

All I know is the null space.

The null space has vectors that solve Ax equals zero.

So this is a guy x. x is in the null space.

Then Ax is zero.

So why is it orthogonal to the rows of A?

If I write down Ax equals zero, which is all I know about the null space, then I guess I want you to see that that's telling me, just that equation right there is telling me that the rows of A, let me write it out.

There's row one of A.

Row two.

Row m of A. that's A.

And it's multiplying X.

And it's producing zero.

OK. Written out that way you'll see it.

So I'm saying that a vector in the row space is perpendicular to this guy X in the null space.

And you see why?

Because this equation is telling you that row one of A multiplying that's a dot product, right?

Row one of A dot product with this x is producing this zero.

So x is orthogonal to the first row.

And to the second row.

Row two of A, x is giving that zero.

Row m of A times x is giving that zero.

So x is -- the equation is telling me that x is orthogonal to all the rows.

Right, it's just sitting there.

That's all we -- it had to be sitting there because we didn't know anything more about the null space than this.

And now I guess to be totally complete here I'd now check that x is orthogonal to each separate row.

But what else strictly speaking do I have to do?

To show that those subspaces are orthogonal, I have to take this x in the null space and show that it's orthogonal to every vector in the row space, every vector in the row space, so what -- what else is in the row space?

This row is in the row space, that row is in the row space, they're all there, but it's not the whole row space, right, we just have to like remember, what does it mean, what does that word space telling us?
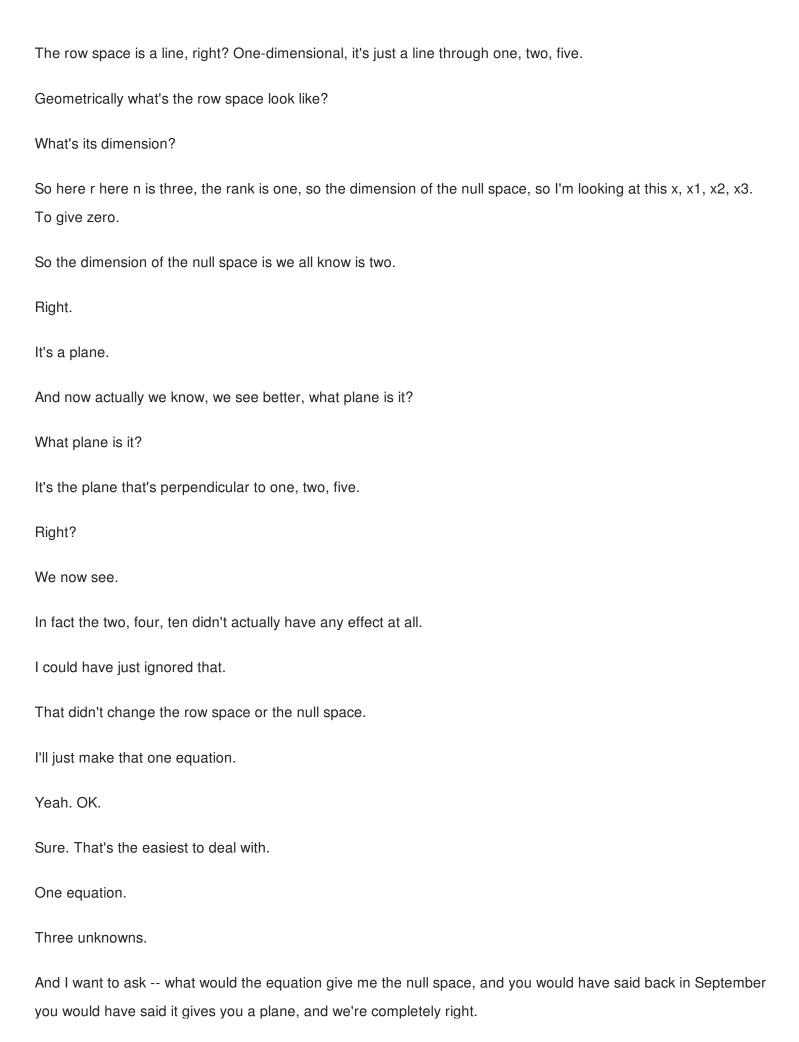
And what else is in the row space?

Besides the rows?

All their combinations.

So I really have to check that sure enough if x is perpendicular to row one, row two, all the different separate rows, then also x is perpendicular to a combination of the rows.

And that's just matrix multiplication again.

You know, I have row one transpose x is zero, so on, row two transpose x is zero, so I'm entitled to multiply that by some c1, this by some c2, I still have zeroes, I'm entitled to add, so I have c1 row one so -- so all this when I put that together that's big parentheses c1 row one plus c2 row two and so on.

Transpose x is zero.

Right?

I just added the zeroes and got zero, and I just added these following the rule.

No big deal.

The whole point was right sitting in that.

OK.

So if I come back to this figure now I'm like a happier person.

Because I have this -- I now see how those subspaces are oriented.

And these subspaces are also oriented.

Well, actually why is that orthogonality?

Well, it's the same statement for A transpose that that one was for A.

So I won't take time to prove it again because we've checked it for every matrix and A transpose is just as good a matrix as A.

So we're orthogonal over there.

So we really have carved up this -- this was like carving up m-dimensional space into two subspaces and this one

was carving up n-dimensional space into two subspaces.

And well, one more thing here.

One more important thing.

Let me move into three dimensions.

Tell me a couple of orthogonal subspaces in three dimensions that somehow don't carve up the whole space, there's stuff left there.

I'm thinking of a couple of orthogonal lines.

If I -- suppose I'm in three dimensions, R^3. And I have one line, one one-dimensional subspace, and a perpendicular one.

Could those be the row space and the null space?

Could those be the row space and the null space?

Could I be in three dimensions and have a row space that's a line and a null space that's a line?

No.

Why not?

Because the dimensions aren't right.

Right? The dimensions are no good.

The dimensions here, r and n-r, they add up to three, they add up to n.

If I take -- just follow that example -- if the row space is one-dimensional, suppose A is what's a good in R^3, I want a one-dimensional row space, let me take one, two, five, two, four, ten.

What's the dimension of that row space?

One.

What's the dimension of the null space?

Tell what's the null space look like in that case?

The row space is a line, right? One-dimensional, it's just a line through one, two, five.

Geometrically what's the row space look like?

What's its dimension?

So here r here n is three, the rank is one, so the dimension of the null space, so I'm looking at this x, x1, x2, x3. To give zero.

So the dimension of the null space is we all know is two.

Right.

It's a plane.

And now actually we know, we see better, what plane is it?

What plane is it?

It's the plane that's perpendicular to one, two, five.

Right?

We now see.

In fact the two, four, ten didn't actually have any effect at all.

I could have just ignored that.

That didn't change the row space or the null space.

I'll just make that one equation.

Yeah. OK.

Sure. That's the easiest to deal with.

One equation.

Three unknowns.

And I want to ask -- what would the equation give me the null space, and you would have said back in September you would have said it gives you a plane, and we're completely right.

And the plane it gives you, the normal vector, you remember in calculus, there was this dumb normal vector called N.

Well there it is.

One, two, five.

OK. What is the what's the point I want to make here?

I want to make -- I want to emphasize that not only are the -- let me write it in words.

So I want to write the null space and the row space are orthogonal, that's this neat fact, which we've -- we've just checked from Ax equals zero, but now I want to say more because there's a little more that's true.

Their dimensions add to the whole space.

So that's like a little extra information.

That it's not like I could have -- I couldn't have a line and a line in three dimensions.

Those don't add up one and one don't add to three.

So I used the word orthogonal complements in R^n.

And the idea of this word complement is that the orthogonal complement of a row space contains not just some vectors that are orthogonal to it, but all.

So what does that mean?

That means that the null space contains all, not just some but all, vectors that are perpendicular to the row space.

OK.

Really what I've done in this half of the lecture is just notice some of the nice geometry that -- that we didn't pick up before because we didn't discuss perpendicular vectors before.

But it was all sitting there.

And now we picked it up.

That these vectors are orthogonal complements.

And I guess I even call this part one of the fundamental theorem of linear algebra.

The fundamental theorem of linear algebra is about these four subspaces, so part one is about their dimension, maybe I should call it part two now.

Their dimensions we got.

Now we're getting their orthogonality, that's part two.

And part three will be about bases for them.

Orthogonal bases.

So that's coming up.

OK.

So I'm happy with that geometry right now.

OK. OK.

Now what's my next goal in this chapter?

Here's the main problem of the chapter.

The main problem of the chapter is -- so this is coming.

It's coming attraction.

This is the very last chapter that's about Ax=b.

I would like to solve that system of equations when there is no solution.

You may say what a ridiculous thing to do.

But I have to say it's done all the time.

In fact it has to be done.

You get -- so the problem is solve -- the best possible solve I'll put quote Ax=b when there is no solution.

And of course what does that mean?

b isn't in the column space.

And it's quite typical if this matrix A is rectangular if I -- maybe I have m equations and that's bigger than the number of unknowns, then for sure the rank is not m, the rank couldn't be m now, so there'll be a lot of right-hand sides with no solution, but here's an example.

Some satellite is buzzing along.

You measure its position.

You make a thousand measurements.

So that gives you a thousand equations for the -- for the parameters that -- that give the position.

But there aren't a thousand parameters, there's just maybe six or something.

Or you're measuring the -- you're doing questionnaires.

You're measuring resistances.

You're taking pulses.

You're measuring somebody's pulse rate.

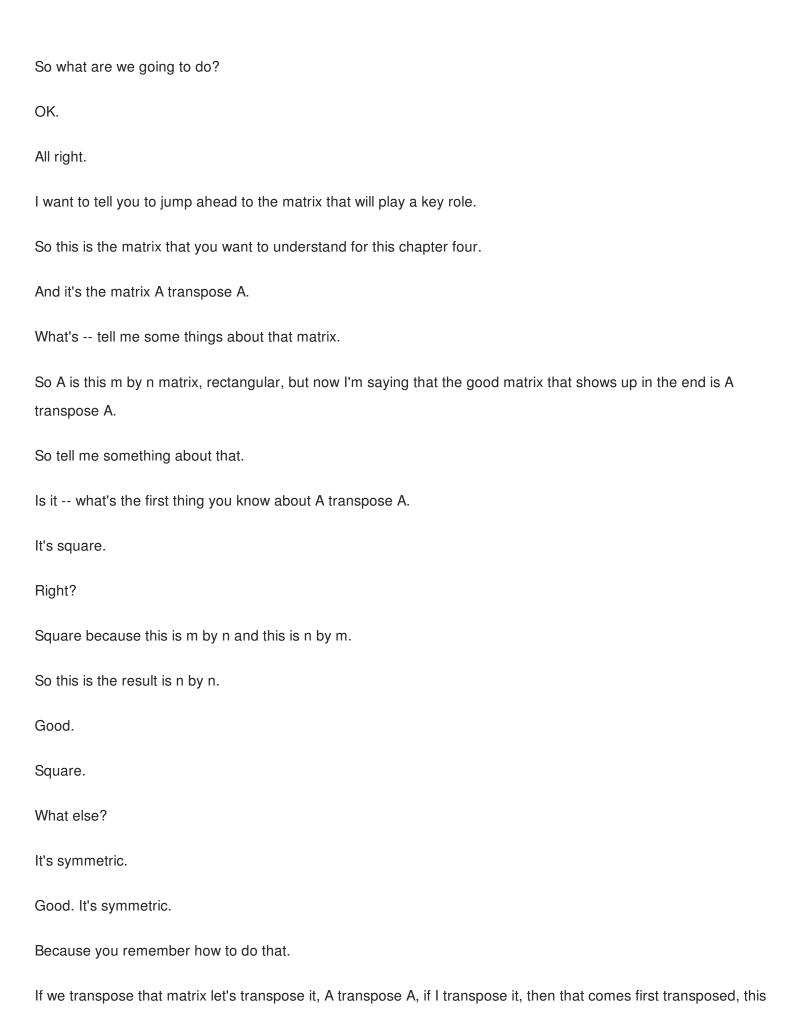There's just one unknown.

OK.

The pulse rate.

So you measure it once, OK, fine, but if you really want to know it, you measure it multiple times, but then the measurements have noise in them, so there's -- the problem is that in many many problems we've got too many equations and they've got noise in the right-hand side.

So Ax=b I can't expect to solve it exactly right, because I don't even know what -- there's a measurement mistake in b.

But there's information too.

There's a lot of information about x in there.

And what I want to do is like separate the noise, the junk, from the information.

And so this is a straightforward linear algebra problem.

How do I solve, what's the best solution?

OK.

Now.

I want to say so that's like describes the problem in an algebraic way.

I got some equations, I'm looking for the best solution.

Well, one way to find it is -- one way to start, one way to find a solution is throw away equations until you've got a nice, square invertible system and solve that.

That's not satisfactory.

There's no reason in these measurements to say these measurements are perfect and these measurements are useless.

We want to use all the measurements to get the best information, to get the maximum information.

But how?

OK. Let me anticipate a matrix that's going to show up.

This A is typically rectangular.

But a matrix that shows up whenever you have -- and we chapter three was all about rectangular matrices.

And we know when this is solvable, you could do elimination on it, right?

But I'm thinking hey, you do elimination and you get equation zero equal other non-zeroes.

I'm thinking we really -- elimination is going to fail.

So that's our question.

Elimination will get us down to -- will tell us if there is a solution or not.

But I'm now thinking not.

So what are we going to do?

OK.

All right.

I want to tell you to jump ahead to the matrix that will play a key role.

So this is the matrix that you want to understand for this chapter four.

And it's the matrix A transpose A.

What's -- tell me some things about that matrix.

So A is this m by n matrix, rectangular, but now I'm saying that the good matrix that shows up in the end is A transpose A.

So tell me something about that.

Is it -- what's the first thing you know about A transpose A.

It's square.

Right?

Square because this is m by n and this is n by m.

So this is the result is n by n.

Good.

Square.

What else?

It's symmetric.

Good. It's symmetric.

Because you remember how to do that.

If we transpose that matrix let's transpose it, A transpose A, if I transpose it, then that comes first transposed, this

comes second, transposed, and then transposing twice is leaves it -- brings it back to the same so it's symmetric.

Good.

Now we now know how to ask more about a matrix.

I'm interested in is it invertible?

If not, what's its null space?

So I want to know about -- because you're going to see, well, let me -- let me even, well I shouldn't do this, but I will.

Let me tell you what equation to solve when you can't solve that one.

The good equation comes from multiplying both sides by A transpose, so the good equation that you get to is this one.

A transpose Ax equals A transpose b.

That will be the central equation in the chapter.

So I think why not tell it to you.

Why not admit it right away.

OK.

I have to -- I should really give x.

I want to sort of indicate that this x isn't I mean this x was the solution to that equation if it existed, but probably didn't.

Now let me give this a different name, x hat.

Because I'm hoping this one will have a solution.

And I'm saying that it's my best solution.

I'll have to say what does best mean.

But that's going to be my -- my plan.

I'm going to say that the best solution solves this equation.

So you see right away why I'm so interested in this matrix A transpose A.

And in its invertibility.

OK. Now, when is it invertible?

OK. Let me take a case, let me just do an example and then -- I'll just pick a matrix here.

Just so we see what A transpose A looks like.

So let me take a matrix A one, one, one, one, two, five.

Just to invent a matrix.

So there's a matrix A.

Notice that it has M equal three rows and N equal two columns.

Its rank is -- the rank of that matrix is two.

Right, yeah, the columns are independent.

Does Ax equal b?

If I look at Ax=b, so x is just x1 x2, and b is b1 b2 b3. Do I expect to solve Ax=b?

What's -- no way, right?

I mean linear algebra's great, but solving three equations with only two unknowns usually we can't do it.

We can only solve it if this vector is b is what?

I can solve that equation if that vector b1 b2 b3 is in the column space.

If it's a combination of those columns then fine.

But usually it won't be.

The combinations just fill up a plane and most vectors aren't on that plane.

So what I'm saying is that I'm going to work with the matrix A transpose A.

And I just want to figure out in this example what A transpose A is.

So it's two by two.

The first entry is a three, the next entry is an eight, this entry is -- what's that entry?

Eight, for sure.

We knew it had to be, and this entry is, what's that now, getting out my trusty calculator, thirty, is that right?

And is that matrix invertible?

Thirty.

There's an A transpose A.

And it is invertible, right?

Three, eight is not a multiple of eight, thirty -- and it's invertible.

And that's the normal, that's what I expect.

So this is I want to show.

So here's the final -- here's the key point.

The null space of A transpose A -- it's not going to be always invertible.

Tell me a matrix -- I have to say that I can't say A transpose A is always invertible.

Because that's asking too much.

I mean what could the matrix A be, for example, so that A transpose A was not invertible?

Well, it even could be the zero matrix.

I mean that's like extreme case.

Suppose I make this rank -- suppose I change to that A.

Now I figure out A transpose A again and I get -- what do I get?

I get nine, I get nine of course and here I get what's that entry?

Twenty-seven. And is that matrix invertible?

No.

And why do I -- I knew it wouldn't be invertible anyway.

Because this matrix only has rank one.

And if I have a product of matrices of rank one, the product is not going to have a rank bigger than one.

So I'm not surprised that the answer only has rank one.

And that's what I -- always happens, that the rank of A transpose A comes out to equal the rank of A.

So, yes, so the null space of A transpose A equals the null space of A, the rank of A transpose A equals the rank of A.

So let's -- as soon as I can why that's true.

But let's draw from that what the fact that I want.

This tells me that this square symmetric matrix is invertible if -- so here's my conclusion.

A transpose A is invertible if exactly when -- exactly if this null space is only got the zero vector.

Which means the columns of A are independent.

So A transpose A is invertible exactly if A has independent columns.

That's the fact that I need about A transpose A.

And then you'll see next time how A transpose A enters everything.

Next lecture is actually a crucial one.

Here I'm preparing for it by getting us thinking about A transpose A.

And its rank is the same as the rank of A, and we can decide when it's invertible.

OK.

So I'll see you Friday.

Thanks.