18.06 Linear Algebra, Fall 1999
Transcript – Lecture 15

OK, guys the -- we're almost ready to make this lecture immortal. OK.

Are we on? All right.

This is an important lecture. It's about projections.

Let me start by just projecting a vector b down on a vector a.

So just to so you see what the geometry looks like in when I'm in -- in just two dimensions, I'd like to find the point along this line so that line through a is a one-dimensional subspace, so I'm starting with one dimension. I'd like to find the point on that line closest to a. Can I just take that problem first and then I'll explain why I want to do it and why I want to project on other subspaces. So where's the point closest to b that's on that line? It's somewhere there. And let me connect that and -- and what's the whole point of my picture now? What's the -- where does orthogonality come into this picture? The whole point is that this best point, that's the projection, P, of b onto the line, where's orthogonality? It's the fact that that's a right angle. That this -- the error -- this is like how much I'm wrong by -- this is the difference between b and P, the whole point is -- that's perpendicular to a.

That's got to give us the equation.

That's got to tell us -- that's the one fact we know, that's got to tell us where that projection is.

Let me also say, look -- I've drawn a triangle there. So if we were doing trigonometry we would do like we would have angles theta and distances that would involve sine theta and cos theta that leads to lousy formulas compared to linear algebra.

The formula that we want comes out nicely and what's the -- what do we know? We know that P, this projection, is some multiple of a, right? It's on that line.

So we know it's in that one-dimensional subspace, it's some multiple, let me call that multiple x, of a. So really it's that number x I'd like to find. So this is going to be simple in 1-D, so let's just carry it through, and then see how it goes in high dimensions. OK.

The key fact is -- the key to everything is that perpendicular. The fact that a is perpendicular to a is perpendicular to e.

Which is (b-ax), xa.

I don't care what -- xa. That that equals zero.

Do you see that as the central equation, that's saying that this a is perpendicular to this -- correction, that's going to tell us what x is.

Let me just raise the board and simplify that and out will come x. OK.

So if I simplify that, let's see, I'll move one to -- one term to one side, the other term will be on the other side, it looks to me like x times a transpose a is equal to a transpose b. Right? I have a transpose b as one f- one term, a transpose a as the other, so right away here's my a transpose a.

But it's just a number now. And I divide by it.

And I get the answer. x is a transpose b over a transpose a. And P, the projection I wanted, is -- that's the right multiple.

That's got a cosine theta built in.

But we don't need to look at angles.

It's -- we've just got vectors here.

And the projection is P is a times that x.

Or x times that a. But I'm really going to -- eventually I'm going to want that x coming on the right-hand side. So do you see that I've got two of the three formulas already, right here, I've got the -- that's the equation -- that leads me to the answer, here's the answer for x, and here's the projection.

OK. can I do add just one more thing to this one-dimensional problem? One more like lift it up into linear algebra, into matrices. Here's the last thing I want to do -- but don't forget those formulas.

a transpose b over a transpose a.

Actually let's look at that for a moment first.

Suppose -- Let me take this next step.

So P is a times x. So can I write that then? P is a times this neat number, a transpose b over a transpose a. That's our projection. Can I ask a couple of questions about it, just while we look, get that digest that formula. Suppose b is doubled.

Suppose I change b to two b. What happens to the projection? So suppose I instead of that vector b that I drew on the board make it two b, twice as long -- what's the projection now? It's doubled too, right? It's going to be twice as far, if b goes twice as far, the projection will go twice as far. And you see it there.

If I put in an extra factor two, then P's got that factor too. Now what about if I double a? What if I double the vector a that I'm projecting onto? What changes? The projection doesn't change at all. Right? Because I'm just -- the line didn't change.

If I double a or I take minus a.

It's still that same line. The projection's still in the same place. And of course if I double a I get a four up above, and I get a four -- an extra four below, they cancel out, and the projection is the same. OK.

So really, this -- I want to look at this as the projection -- there's a matrix here. The projection is carried out by some matrix that I'm going to call the projection matrix. And in other words the projection is some matrix that acts on this guy b and produces the projection. The projection P is the projection matrix acting on whatever the input is. The input is b, the projection matrix is P. OK.

Actually you can tell me right away what this projection matrix is. So this is a pretty interesting matrix. What matrix is multiplying b? I'm just -- just from my formula -- I see what P is.

P, this projection matrix, is -- what is it? I see a a transpose above, and I see a transpose a below.

And those don't cancel. That's not one.

Right? That's a matrix.

Because down here, the a transpose a, that's just a number, a transpose a, that's the length of a squared, and up above is a column times a row. Column times a row is a matrix. So this is a full-scale n by n matrix, if I -- if I'm in n dimensions. And it's kind of an interesting one. And it's the one which if I multiply by b then I get this, you see once again I'm putting parentheses in different places. I'm putting the parentheses right there. I'm saying OK, that's really the matrix that produces this projection.

OK. Now, tell me -- all right, what are the properties of that matrix? I'm just using letters here, a and b, I could put in numbers, but I think it's -- for once, it's clearer with letters, because all formulas are simple, a transpose b over a transpose a -- that's the number that multiplies the a, and then I see wait a minute, there's a matrix here and what's the rank of that matrix, by the way? What's the rank of that matrix, yeah -- let me just ask you about that matrix.

Which looks a little strange, a a transpose over this number.

But well, I could ask you its column space.

Yeah, let me ask you its column space.

So what's the column space of a matrix? If you multiply that matrix by anything you always get in the column space, right? The column space of a matrix is when you multiply any vector by that matrix -- any vector b, by the matrix, you always land in the column space. That's what column spaces work.

Now what space do we always land in? What's the column space of -- what's the result when I multiply this any vector b by my matrix? So I have P times b, where I? I'm on that line, right? The column space, so here are facts about this matrix. The column space of P, of this projection matrix, is the line through a.

And the rank of this matrix is you can all say it at once one.

Right. The rank is one. This is a rank one matrix. Actually it's exactly the form that we're familiar with a rank one matrix.

A column times a row, that's a rank one matrix, that column is the basis for the column space.

Just one dimension. OK.

So I know that much about the matrix.

But now there are two more facts about the matrix that I want to notice. First of all is the matrix symmetric? That's a natural question for matrices. And the answer is yes.

If I take the transpose of this -- there's a number down there, the transpose of a a transpose is a a transpose. So P is symmetric.

P transpose equals P. So this is a key property.

That the projection matrix is symmetric.

One more property now and this is the real one.

What happens if I do the projection twice? So I'm looking for something, some information about P squared. But just give me in terms of that picture, in terms my picture, I take any vector b, I multiply it by my projection matrix, and that puts me there, so this is Pb.

And now I project again. What happens now? What happens when I apply the projection matrix a second time? To this, so I'm applying it once brings me here and the second time brings me I stay put.

Right? The projection for a point on this line the projection is right where it is.

The projection is the same point.

So that means that if I project twice, I get the same answer as I did in the first projection. So those are the two properties that tell me I'm looking at a projection matrix.

It's symmetric and it's square is itself.

Because if I project a second time, it's the same result as the first result. OK. So that's -- and then here's the exact formula when I know what I'm projecting onto, that line through a, then I know what P is. So do you see that I have all the pieces here for projection on a line? Now, and those -- please remember those.

So there are three formulas to remember.

The formula for x, the formula for P, which is just ax, and then the formula for capital P, which is the matrix. Good.

Good. OK.

Now I want to move to more dimensions.

So we're going to have three formulas again but you'll have to -- they'll be a little different because we won't have a single line but -- a plane or three-dimensional or a n-dimensional subspace.

OK. So now I'll move to the next question. Maybe -- let me say first why I want this projection, and then we'll figure out what it is, we'll go completely in parallel there, and then we'll use it. OK, why do I want this projection? Well, so why project? It's because I'm as I mentioned last time this new chapter deals with equations Ax=b may have no solution.

So that's really my problem, that I'm given a bunch of equations probably too many equations, more equations than unknowns, and I can't solve them. OK. So what do I do? I solve the closest problem that I can solve.

And what's the closest one? Well, ax will always be in the column space of a. That's my problem.

My problem is ax has to be in the column space and b is probably not in the column space.

So I change b to what? I choose the closest vector in the column space, so I'll solve Ax equal P instead. That one I can do.

Where P is this is the projection of b onto the column space. That's why I want to be able to do this. Because I have to find a solution here, and I'm going to put a little hat there to indicate that it's not the x, it's not the x that doesn't exist, it's the x hat that's best possible. So I must be able to figure out what's the good projection there.

What's the good right-hand side that is in the column space that's as close as possible to b and then I'm -- then I know what to do. OK.

So now I've got that problem. So that's why I have the problem again but now let me say I'm in three dimensions, so I have a plane maybe for example, and I have a vector b that's not in the plane. And I want to project b down into the plane. OK.

So there's my question. How do I project a vector and I'm -- what I'm looking for is a nice formula, and I'm counting on linear algebra to just come out right, a nice formula for the projection of b into the plane.

The nearest point. So this again a right angle is going to be crucial. OK.

Now so what's -- first of all I've got to say what is that plane. To get a formula I have to tell you what the plane is. How I going to tell you a plane? I'll tell you a basis for the plane, I'll tell you two vectors a one and a two that give you a basis for the plane, so that -- let us say -- say there's an a one and here's an a -- a vector a two.

They don't have to be perpendicular.

But they better be independent, because then that tells me the plane. The plane is the -- is the plane of a one and a two. And actually going back to my -- to this connection, this plane is a column space, it's the column space of what matrix? What matrix, so how do I connect the two questions? I'm thinking how do I project onto a plane and I want to get a matrix in here.

Everything's cleaner if I write it in terms of a matrix.

So what matrix has these -- has that column space? Well of course it's just the matrix that has a one in the first column and a two in the second column.

Right, just just let's be sure we've got the question before we get to the answer. So I'm looking for again I'm given a matrix a with two columns.

And really I'm ready once I get to two I'm ready for n.

So it could be two columns, it could be n columns.

I'll write the answer in terms of the matrix a.

And the point will be those two columns describe the plane, they describe the column space, and I want to project.

OK. And I'm given a vector b that's probably not in the column space.

Of course, if b is in the column space, my projection is simple, it's just b.

But most likely I have an error e, this b minus P part, which is probably not zero. OK. But the beauty is that I know -- from geometry or I could get it from calculus or I could get it from linear algebra that that this this vector -- this is the part of b that's that's perpendicular to the plane.

That e is perpendicular is perpendicular to the plane.

If your intuition is saying that that's the crucial fact.

That's going to give us the answer.

OK. So let me, that's the problem.

Now for the answer. So this is a lecture that's really like moving along. Because I'm just plotting that problem up there and asking you what combination -- now, yeah, so what is it? P.

What is this projection P? This is projection P, is some multiple of these basis guys, right, some multiple of the columns. But I don't like writing out x one a one plus x two a two, I would rather right that as ax. Well, actually I should put if I'm really doing everything right, I should put a little hat on it -- to remember that this x -- that those are the numbers and I could have a put a hat way back there is right, so this is this is the projection, P. P is ax bar.

And I'm looking for x bar. So that's what I want an equation for. So now I've got hold of the problem. The problem is find the right combination of the columns so that

the error vector is perpendicular to the plane. Now let me turn that into an equation. So I'll raise the board and just turn that -- what we've just done into an equation.

So let me I'll write again the main point.

The projection is ax b- x hat. And our problem is find x hat.

And the key is that b minus ax hat, that's the error.

This is the e. Is perpendicular to the plane.

That's got to give me well what I looking for, I'm looking for two equations now because I've got an x one and an x two. And I'll get two equations because so this thing e is perpendicular to the plane.

So what does that mean? I guess it means it's perpendicular to a one and also to a two.

Right, those are two vectors in the plane and the things that are perpendicular to the plane are perpendicular to a one and a two. Let me just repeat.

This this guy then is perpendicular to the plane so it's perpendicular to that vector and that vector.

Not -- it's perpendicular to that of course.

But it's perpendicular to everything I the plane.

And the plane is really told me by a one and a two.

So really I have the equations a one transpose b minus ax is zero. And also a two transpose b minus ax is zero. Those are my two equations.

But I want those in matrix form.

I want to put those two equations together as a matrix equation and it's just comes out right.

Look at the matrix a transpose. Put a one a one transpose is its first row, a two transpose is its second row, that multiplies this b-ax, and gives me the zero and the zero. I'm you see the -- this is one way -- to come up with this equation. So the equation I'm coming up with is a transpose b-ax hat is zero. OK.

That's my equation. All right.

Now I want to stop for a moment before I solve it and just think about it. First of all do you see that that equation back in the very first problem I solved on a line, what was -- what was on a line the matrix a only had one column, it was just little a.

So in the first problem I solved, projecting on a line, this for capital a you just change that to little a and you have the same equation that we solved before. a transpose e equals zero. OK.

Now a second thing, second comment.

I would like to since I know about these four subspaces, I would like to get them into this picture.

So let me ask the question, what subspace is this thing in? Which of the four subspaces is that error vector e, this is this is nothing but e -- this is this guy, coming in down perpendicular to the plane. What subspace is e in? From this equation. Well the equation is saying a transpose e is zero. So I'm learning here that e is in the null space of a transpose.

Right? That's my equation.

And now I just want to see hey of course that that was right.

Because things that are in the null space of a transpose, what do we know about the null space of a transpose? So that last lecture gave us the sort of the geometry of these subspaces. And the orthogonality of them.

And do you remember what it was? What on the right side of our big figure we always have the null space of a transpose and the column space of a, and they're orthogonal. So e in the null space of a transpose is saying e is perpendicular to the column space of a. Yes.

I just feel OK, the damn thing came out right.

The equation for the equation that I struggled to find for e really said what I wanted, that the error e is perpendicular to the column space of a, just right.

And from our four fundamental subspaces we knew that that is the same as that. To say e is in the null space of a transpose says e's perpendicular to the column space. OK.

So we've got this equation. Now let's just solve it. All right. Let me just rewrite it as a transpose a x hat equals a transpose b.

That's our equation. That gives us x.

And -- allow me to keep remembering the one-dimensional case. The one-dimensional case, this was little a. So this was just a number, little a transpose, a transpose a was just a vector row times a column, a number.

And this was a number. And x was the ratio of those numbers. But now we've got matrices, this one is n by n. a transpose a is an n by n matrix. OK.

So can I move to the next board for the solution? OK. This is the -- the key equation. Now I'm ready for the formulas that we have to remember. What's x hat? What's the projection, what's the projection matrix, those are my three questions. That we answered in the 1-D case and now we're ready for in the n-dimensional case.

So what is x hat? Well, what can I say but a transpose a inverse, a transpose b. OK. That's the solution to -- to our equation. What's the projection? That's more interesting. What's the projection? The projection is a x hat. That's how x hat got

into the picture in the first place. x hat was the was the combination of columns in the I had to look for those numbers and now I found them. Was the combination of the columns of a that gave me the projection.

OK. So now I know what this guy is.

So it's just I multiply by a. a a transpose a inverse a transpose b. That's looking a little messy but it's not bad. That that combination is is our like magic combination. This is the thing which is which use which is like what's it like, what was it in one dimension? What was that we had this we must have had this thing way back at the beginning of the lecture. What did we -- oh that a was just a column so it was little a, little a transpose over a transpose a, right, that's what it was in 1-D.

You see what's happened in more dimensions, I -- I'm not allowed to to just divide because because I don't have a number, I have to put inverse, because I have an n by n matrix.

But same formula. And now tell me what's the projection matrix? What matrix is multiplying b to give the projection? Right there.

Because there it -- I even already underlined it by accident. The projection matrix which I use capital P is this, it's it's that thing, shall I write it again, a times a transpose a inverse times a transpose. Now if you'll bear with me I'll think about what have I done here.

I've got this formula. Now the first thing that occurs to me is something bad. Look why don't I just you know here's a product of two matrices and I want its inverse, why don't I just use the formula I know for the inverse of a product and say OK, that's a inverse times a transpose inverse, what will happen if I do that? What will happen if I say hey this is a inverse times a transpose inverse, then shall I do it? It's going to go on videotape if I do it, and I don't -- all right, I'll put it there, but just like don't take the videotape quite so carefully. OK.

So if I put that thing it -- it would be a a inverse a transpose inverse a transpose and what's that? That's the identity. But what's going on? So why -- you see my question is somehow I did something wrong. That that wasn't allowed.

And and and why is that? Because a is not a square matrix. a is not a square matrix. It doesn't have an inverse. So I have to leave it that way.

This is not OK. If if a was a square invertible matrix, then then I couldn't complain.

Yeah I think -- let me think about that case.

But you but my main case, the whole reason I'm doing all this, is that a is this matrix that has x too many rows, it's just got a couple of columns, like a one and a two, but lots of rows. Not square.

And if it's not square, this a transpose a is square but I can't pull it apart like this -- I'm not allowed to do this pull apart, except if a was square.

Now if a is square what's what's going on if a is a square matrix? a nice square inv- invertible matrix. Think.

What's up with that what's with that case.

So this is that the formula ought to work then too.

If a is a nice square invertible matrix what's its column space, so it's a nice n by n invertible everything great matrix, what's its column space, the whole of R^n. So what's the projection matrix if I'm projecting onto the whole space? It's the identity matrix right? If I'm projecting b onto the whole space, not just onto a plane, but onto all of 3-D, then b is already in the column space, the projection is the identity, and this is gives me the correct formula, P is I. But if I'm projecting onto a subspace then I can't split those apart and I have to stay with that formula. OK.

And what can I say if -- so I remember this formula for 1-D and that's what it looks like in n dimensions.

And what are the properties that I expected for any projection matrix? And I still expect for this one? That matrix should be symmetric and it is. P transpose is P. Because if I transpose this, this guy's symmetric, and its inverse is symmetric, and if I transpose this one when I transpose it will pop up there, become a, that a transpose will pop up here, and I'm back to P again.

And do we dare try the other property which is P squared equal P? It's got to be right.

Because we know geometrically that the first projection pops us onto the column space and the second one leaves us where we are. So I expect that if I multiply by let me do it -- if I multiply by another P, so there's another a, another a transpose a inverse a transpose, can you -- eight (a)-s in a row is quite obscene but -- do you see that it works? So I'm squaring that so what do I do-- how do I see that multiplication? Well, yeah, I just want to put parentheses in good places, so I see what's happening, yeah, here's an a transpose a sitting together -- so when that a transpose a multiplies its inverse, all that stuff goes, right.

And leaves just the a transpose at the end, which is just what we want. So P squared equals P.

So sure enough those two properties hold.

OK. OK we really have got now all the formulas. x hat, the projection P, and the projection matrix capital P.

And now my job is to use them. OK. So when would I have a bunch of equations, too many equations and yet I want the best answer and the -- the most important example, the most common example is if I have points so here's the -- here's the application.

v squared. Fitting by a line.

OK. So I'll start this application today and there's more in it than I can do in this same lecture. So that'll give me a chance to recap the formulas and there they are, and recap the ideas.

So let me start the problem today.

I'm given a bunch of data points.

And they lie close to a line but not on a line.

Let me take that. Say a t equal to one, two and three, I have one, and two and two again. So my data points are this is the like the time direction and this is like well let me call that b or y or something. I'm given these three points and I want to fit them by a line. By the best straight line. So the problem is fit the points one, one is the first point -- the second point is t equals two, b equal one, and the third point is t equal three, b equal to two.

So those are my three points, t equal sorry,that's two.

Yeah, OK. So this is the point one, one. This is the point two, two, and that's the point three, two.

And of course there isn't a -- a line that goes through them.

So I'm looking for the best line.

I'm looking for a line that probably goes somewhere, do you think it goes somewhere like that? I didn't mean to make it go through that point, it won't. It'll kind of -- it'll go between so the error there and the error there and the error there are as small as I can get them. OK, what I'd like to do is find the matrix a.

Because once I've found the matrix a the formulas take over.

So what I'm looking for this line, b is C+Dt.

So this is in the homework that I sent out for today.

Find the best line. So I'm looking for these numbers. C and D.

That tell me the line and I want them to be as close to going through those three points as I can get.

I can't get exactly so there are three equations to go through the three points. It would it will go exactly through that point if let's see that first point has t equal to one, so that would say C+D equaled 1.

This is the one, one. The second point t is two. So C+2D should come out to equal 2. But I also want to get the third equation in and at that third equation t is three so C+3D equals only 2. That's the key. Is to write down what equations we would like to solve but can't. Reason we if we could solve them that would mean that we could put a line through all three points and of course if these numbers one, two, two were different, we could do it.

But with those numbers, one, two, two, we can't. So what is our equation Ax equal Ax equal b that we can't solve? I just want to say what's the matrix here, what's the unknown x, and what's the right-hand side.

So this is the matrix is one, one, one, one, two, three. The unknown is C and D.

And the right-hand side if one, two, two.

Right, I've just taken my equations and I've said what is Ax and what is b. Then there's no solution, this is the typical case where I have three equations -- two unknowns, no solution, but I'm still looking for the best solution.

And the best solution is taken is is to solve not this equation Ax equal b which has which has no solution but the equation that does have a solution, which was this one.

So that's the equation to solve.

That's the central equation of the subject.

I can't solve Ax=b but magically when I multiply both sides by a transpose I get an equation that I can solve and its solution gives me x, the best x, the best projection, and I discover what's the matrix that's behind it. OK.

So next time I'll complete an example, numerical example.

today was all letters, numbers next time.

Thanks.

MIT OpenCourseWare
http://ocw.mit.edu

18.06SC Linear Algebra
Fall 2011