

MIT OpenCourseWare
<http://ocw.mit.edu>

18.06 Linear Algebra, Spring 2005

Please use the following citation format:

Gilbert Strang, *18.06 Linear Algebra, Spring 2005*. (Massachusetts Institute of Technology: MIT OpenCourseWare). <http://ocw.mit.edu> (accessed MM DD, YYYY). License: Creative Commons Attribution-Noncommercial-Share Alike.

Note: Please use the actual date you accessed this material in your citation.

For more information about citing these materials or our Terms of Use, visit:
<http://ocw.mit.edu/terms>

MIT OpenCourseWare
<http://ocw.mit.edu>

18.06 Linear Algebra, Spring 2005
Transcript – Lecture 32

OK, here we go with, quiz review for the third quiz that's coming on Friday. So, one key point is that the quiz covers through chapter six. Chapter seven on linear transformations will appear on the final exam, but not on the quiz. So I won't review linear transformations today, but they'll come into the full course review on the very last lecture.

So today, I'm reviewing chapter six, and I'm going to take some old exams, and I'm always ready to answer questions.

And I thought, kind of help our memories if I write down the main topics in chapter six.

So, already, on the previous quiz, we knew how to find eigenvalues and eigenvectors. Well, we knew how to find them by that determinant of $A - \lambda I$ equals zero. But, of course, there could be shortcuts. There could be, like, useful information about the eigenvalues that we can speed things up with. OK.

Then, the new stuff starts out with a differential equation, so I'll do a problem. I'll do a differential equation problem first. What's special about symmetric matrices? Can we just say that in words? I'd better write it down, though.

What's special about symmetric matrices? Their eigenvalues are real. The eigenvalues of a symmetric matrix always come out real, and there always are enough eigenvectors. Even if there are repeated eigenvalues, there are enough eigenvectors, and we can choose those eigenvectors to be orthogonal.

So if A equals A transposed, the big fact will be that we can diagonalize it, and those eigenvector matrix, with the eigenvectors in the column, can be an orthogonal matrix. So we get a $Q \lambda Q^T$. That, in three symbols, expresses a wonderful fact, a fundamental fact for symmetric matrices. OK.

Then, we went beyond that fact to ask about positive definite matrices, when the eigenvalues were positive.

I'll do an example of that. Now we've left symmetry.

Similar matrices are any square matrices, but two matrices are similar if they're related that way.

And what's the key point about similar matrices? Somehow, those matrices are representing the same thing in different basis, in chapter seven language.

In chapter six language, what's up with these similar matrices? What's the key fact, the key positive fact about similar matrices? They have the same eigenvalues. Same eigenvalues.

So if one of them grows, the other one grows.

If one of them decays to zero, the other one decays to zero.

Powers of A will look like powers of B, because powers of A and powers of B only differ by an M inverse and an M way on the outside. So if these are similar, then B to the k-th power is M inverse A to the k-th power M.

And that's why I say, eh, this M, it does change the eigenvectors, but it doesn't change the eigenvalues. So same lambdas. And then, finally, I've got to review the point about the SVD, the Singular Value Decomposition. OK.

So that's what this quiz has got to cover, and now I'll just take problems from earlier exams, starting with a differential equation.

OK. And always ready for questions.

So here is an exam from about the year zero, and it has a three by three. So that was -- but it's a pretty special-looking matrix, it's got zeroes on the diagonal, it's got minus ones above, and it's got plus ones like that. So that's the matrix A.

OK. Step one is, well, I want to solve that equation.

I want to find the general solution.

I haven't given you a $u(0)$ here, so I'm looking for the general solution, so now what's the form of the general solution? With three arbitrary constants going to be inside it, because those will be used to match the initial condition. So the general form is u at time t is some multiple of the first special solution.

The first special solution will be growing like the eigenvalue, and it's the eigenvector. So that's a pure exponential solution, just staying with that eigenvector.

Of course, I haven't found, yet, the eigenvalues and eigenvectors. That's, normally, the first job. Now, there will be second one, growing like e to the lambda two, and a third one growing like e to the lambda three. So we're all done -- well, we haven't done anything yet, actually.

I've got to find the eigenvalues and eigenvectors, and then I would match $u(0)$ by choosing the right three constants. OK.

So now I ask -- ask you about the eigenvalues and eigenvectors, and you look at this matrix and what do you see in that matrix? Um, well, I guess we might ask ourselves right away, is it singular? Is it singular? Because, if so, then we really have a head start, we know one of the eigenvalues is zero. Is that matrix singular? Eh, I don't know, do you take the determinant to find out? Or maybe you look at the first row and third row and say, hey, the first row and third row are just opposite signs, they're linear-dependent? The first column and third column are dependent -- it's singular. So one eigenvalue is zero. Let's make that lambda one. Lambda one, then, will be zero. OK.

Now we've got a couple of other eigenvalues to find, and, I suppose the simplest way is to look at $A - \lambda I$. So let me just put minus λ in here, minus ones above, ones below. But, actually, before I do it, that matrix is not symmetric, for sure, right? In fact, it's the very opposite of symmetric. That matrix A transpose, how is A transpose connected to A ? It's negative A . It's an anti-symmetric matrix, skew-symmetric matrix. And we've met, maybe, a two-by-two example of skew-symmetric matrices, and let me just say, what's the deal with their eigenvalues? They're pure imaginary.

They'll be on the imaginary axis, there be some multiple of I if it's an anti-symmetric, skew-symmetric matrix.

So I'm looking for multiples of I , and of course, that's zero times I , that's on the imaginary axis, but maybe I just do it out, here.

λ^3 . well, maybe that's minus λ^3 , and then a zero and a zero.

Zero, and then maybe I have a plus a λ , and another plus λ , but those go with a minus sign.

Am I getting minus two λ equals zero? So. So I'm solving $\lambda^3 + 2\lambda = 0$. So one root factors out λ , and the the rest is $\lambda^2 + 2$.

OK. This is going the way we expect, right? Because this gives the root $\lambda = 0$, and gives the other two roots, which are $\lambda = \pm \sqrt{2}i$. The solutions of when is $\lambda^2 + 2 = 0$ then the eigenvalues those guys, what are they? They're a multiple of i , they're just square root of two i .

When I set this equals to zero, I have $\lambda^2 = -2$, right? To make that zero? And the roots are square root of two i and minus the square root of two i . So now I know what those are.

I'll put those in, now.

Either the zero t is just a one.

That's just a one. This is square root of two i and this is minus square root of two i .

So, is the solution decaying to zero? Is this a completely stable problem where the solution is going to zero? No.

In fact, all these things are staying the same size.

This thing is getting multiplied by this number.

$e^{i\omega t}$ something t , that's a number that has magnitude one, and sort of wanders around the unit circle. Same for this.

So that the solution doesn't blow up, and it doesn't go to zero. OK.

And to find out what it actually is, we would have to plug in initial conditions. But actually, the next question I ask is, when does the solution return to its initial value?

I won't even say what's the initial value. This is a case in which I think this solution is periodic after. At t equals zero, it starts with c_1 , c_2 , and c_3 , and then at some value of t , it comes back to that.

So that's a very special question, Well, let's just take three seconds, because that special question isn't likely to be on the quiz. But it comes back to the start, when? Well, whenever we have e to the two πi , that's one, and we've come back again.

So it comes back to the start. It's periodic, when this square root of two i -- shall I call it capital T , for the period? For that particular T , if that equals two πi , then e to this thing is one, and we've come around again. So the period is T is determined here, cancel the i -s, and T is π times the square root of two.

So that's pretty neat. We get all the information about all solutions, we haven't fixed on only one particular solution, but it comes around again.

So this was probably my first chance to say something about the whole family of anti-symmetric, skew-symmetric matrices. OK.

And then, finally, I asked, take two eigenvectors (again, I haven't computed the eigenvectors) and it turns out they're orthogonal. They're orthogonal.

The eigenvectors of a symmetric matrix, or a skew-symmetric matrix, are always orthogonal. I guess my conscience makes me tell you, what are all the matrices that have orthogonal eigenvectors? And symmetric is the most important class, so that's the one we've spoken about. But let me just put that little fact down, here. Orthogonal x -s. eigenvectors. A matrix has orthogonal eigenvectors, the exact condition -- it's quite beautiful that I can tell you exactly when that happens.

It happens when A times A transpose equals A transpose times A . Any time that's the condition for orthogonal eigenvectors. And because we're interested in special families of vectors, tell me some special families that fit. This is the whole requirement.

That's a pretty special requirement most matrices have.

So the average three-by-three matrix has three eigenvectors, but not orthogonal. But if it happens to commute with its transpose, then, wonderfully, the eigenvectors are orthogonal. Now, do you see how symmetric matrices pass this test? Of course. If A transpose equals A , then both sides are A squared, we've got it.

How do anti-symmetric matrices pass this test? If A transpose equals minus A , then we've got it again, because we've got minus A squared on both sides.

So that's another group. And finally, let me ask you about our other favorite family, orthogonal matrices. Do orthogonal matrices pass this test, if A is a Q , do they pass the test for orthogonal eigenvectors. Well, if A is Q , an orthogonal matrix, what is Q transpose Q ? It's I . And what is $Q Q$ transpose? It's I , we're talking square matrices here.

So yes, it passes the test. So the special cases are symmetric, anti-symmetric (I'll say skew-symmetric,) and orthogonal. Those are the three important special classes that are in this family.

OK. That's like a comment that, could have been made back in, section six point four.

OK, I can pursue the differential equations, also this question, didn't ask you to tell me, how would I find this matrix exponential, e^{At} ? So can I erase this? I'll just stay with this same... how would I find e^{At} ? Because, how does that come in? That's the key matrix for a differential equation, because the solution is -- the solution is $u(t)$ is $e^{(At)} u(0)$. So this is like the fundamental matrix that multiplies the given function and gives the answer.

And how would we compute it if we wanted that? We don't always have to find e^{At} , because I can go directly to the answer without any e^{At} -s, but hiding here is an e^{At} , and how would I compute it? Well, if A is diagonalizable. So I'm now going to put in my usual if A can be diagonalized (and everybody remember that there is an if there, because it might not have enough eigenvectors) this example does have enough, random matrices have enough. So if we can diagonalize, then we get a nice formula for this, because an S comes way out at the beginning, and S inverse comes way out at the end, and we only have to take the exponential of λ .

And that's just a diagonal matrix, so that's just $e^{\lambda t}$, these guys are showing up, now, in $e^{\lambda t}$. OK? That's a really quick review of that formula.

It's something we can compute it quickly if we have done the S and λ part. If we know S and λ , then it's not hard to take that step.

OK, that's some comments on differential equations.

I would like to go on to a next question that I started here.

And it's, got several parts, and I can just read it out.

What we're given is a three-by-three matrix, and we're told its eigenvalues, except one of these is, like, we don't know, and we're told the eigenvectors. And I want to ask you about the matrix. OK.

So, first question. Is the matrix diagonalizable? And I really mean for which c , because I don't know c , so my questions will all be, for which is there a condition on c , does one c work. But your answer should tell me all the c -s that work. I'm not asking for you to tell me, well, c equal four, yes, that checks out.

I want to know all the c -s that make it diagonalizable.

OK? What's the real on diagonalizable? We need enough eigenvectors, right? We don't care what those eigenvalues are, it's eigenvectors that count for diagonalizable, and we need three independent ones, and are those three guys independent? Yes.

Actually, let's look at them for a moment.

What do you see about those three vectors right away? They're more than independent. Can you see why those three got chosen? Because it will come up in

the next part, they're orthogonal. Those eigenvectors are orthogonal. They're certainly independent.

So the answer to diagonalizable is, yes, all c , all c . Doesn't matter.

c could be a repeated guy, but we've got enough eigenvectors, so that's what we care about.

OK, second question. For which values of c is it symmetric? OK, what's the answer to that one? If we know the same setup if we know that much about it, we know those eigenvectors, and we've noticed they're orthogonal, then which c -s will work? So the eigenvalues of that symmetric matrix have to be real.

So all real c . If c was i , the matrix wouldn't have been symmetric. But if c is a real number, then we've got real eigenvalues, we've got orthogonal eigenvectors, that matrix is symmetric. OK, positive definite. OK, now this is a sub-case of symmetric, so we need c to be real, so we've got a symmetric matrix, but we also want the thing to be positive definite.

Now, we're looking at eigenvalues, we've got a lot of tests for positive definite, but eigenvalues, if we know them, is certainly a good, quick, clean test. Could this matrix be positive definite? No.

No, because it's got an eigenvalue zero.

It could be positive semi-definite, you know, like consolation prize, if c was greater or equal to zero, it would be positive semi-definite.

But it's not, no.

Semi-definite, if I put that comment in, semi-definite, that the condition would be c greater or equal to zero. That would be all right.

OK. Next part.

Is it a Markov matrix? Hm.

Could this matrix be, if I choose the number c correctly, a Markov matrix? Well, what do we know about Markov matrices? Mainly, we know something about their eigenvalues. One eigenvalue is always one, and the other eigenvalues are smaller.

Not larger. So an eigenvalue two can't happen. So the answer is, no, not a matrix that's never a Markov matrix.

OK? And finally, could one half of A be a projection matrix? So could it- could this -- eh-eh could this be twice a projection matrix? So let me write it this way.

Could A over two be a projection matrix? OK, what are projection matrices? They're real. I mean, they're symmetric, so their eigenvalues are real. But more than that, we know what those eigenvalues have to be.

What do the eigenvalues of a projection matrix have to be? See, that any nice matrix we've got an idea about its eigenvalues. So the eigenvalues of projection matrices are zero and one.

Zero and one, only.

Because $P^2 = P$, let me call this matrix P , so $P^2 = P$, so $\lambda^2 = \lambda$, because eigenvalues of P^2 are λ^2 , and we must have that, so $\lambda = 0$ or 1 .

OK. Now what value of c will work there? So, then, there are some value that will work, and what will work? $c = 0$ will work, or what else will work? $c = 2$. Because if $c = 2$, then when we divide by two, this Eigenvalue of two will drop to one, and so will the other one, so, or $c = 2$. OK, those are the guys that will work, and it was the fact that those eigenvectors were orthogonal, the fact that those eigenvectors were orthogonal carried us a lot of the way, here.

If they weren't orthogonal, then symmetric would have been dead, positive definite would have been dead, projection would have been dead.

But those eigenvectors were orthogonal, so it came down to the eigenvalues. OK, that was like a chance to review a lot of this chapter. Shall I jump to the singular value decomposition, then, as the third, topic for, for the review? OK, so I'm going to.

jump to this. OK.

So this is the singular value decomposition, known to everybody as the SVD. And that's a factorization of A into orthogonal times diagonal times orthogonal.

And we always call those U and σ and V^T .

OK. And the key to that -- this is for every matrix, every A , every A .

Rectangular, doesn't matter, whatever, has this decomposition.

So it's really important. And the key to it is to look at things like $A^T A$. Can we remember what happens with $A^T A$? If I just transpose that I get $V^T \sigma^T U^T$, that's multiplying A , which is $U \sigma V^T$, and the result is V on the outside, $U^T U$ is the identity, because it's an orthogonal matrix. So I'm just left with $\sigma^T \sigma$ in the middle, that's a diagonal, possibly rectangular diagonal by its transpose, so the result, this is orthogonal, diagonal, orthogonal. So, I guess, actually, this is the SVD for $A^T A$.

Here I see orthogonal, diagonal, and orthogonal.

Great. But a little more is happening. For $A^T A$, the orthogonal guys are the same.

It's V and V^T . What I'm seeing here? I'm seeing the factorization for a symmetric matrix.

This thing is symmetric. So in a symmetric case, U is the same as V . U is the same as V for this symmetric matrix, and, of course, we see it happening. OK.

So that tells us, right away, what V is.

V is the eigenvector matrix for A transpose A .

OK. Now, if you were here when I lectured about this topic, when I gave the topic on singular value decompositions, you'll remember that I got into trouble. I'm sorry to remember that myself, but it happened. OK.

How did it happen? I was in great shape for a while, cruising along. So I found the eigenvectors for A transpose A . Good.

I found the singular values, what were they? What were the singular values? The singular value number i , or -- these are the guys in σ -- this is diagonal with the number σ in it. This diagonal is σ one, σ two, up to the rank, σ r , those are the non-zero ones. So I found those, and what are they? Remind me about that? Well, here, I'm seeing them squared, so their squares are the eigenvalues of A transpose A .

Good. So I just take the square root, if I want the eigenvalues of A transpose -- If I want the σ s and I know these, I take the square root, the positive square root. OK.

Where did I run into trouble? Well, then, my final step was to find U . And I didn't read the book.

So, I did something that was practically right, but -- well, I guess practically right is not quite the same. OK, so I thought, OK, I'll look at A transpose A . What happened when I looked at A transpose A ? Let me just put it here, and then I can feel it. OK, so here's A transpose A . So that's U σ V transpose, that's A , and then the transpose is V σ transpose, U σ transpose. Fine. And then, in the middle is the identity again, so it looks great.

U σ σ transpose, U transpose.

Fine. All good, and now these columns of U are the eigenvectors, that's U is the eigenvector matrix for this guy. That was correct, so I did that fine. Where did something go wrong? A sign went wrong. A sign went wrong because -- and now -- now I see, actually, somebody told me right after class, we can't tell from this description which sign to give the eigenvectors. If these are the eigenvectors of this matrix, well, if you give me an eigenvector and I change all its signs, we've still got another eigenvector.

So what I wasn't able to determine (and I had a fifty-fifty chance and life let me down,) the signs I just happened to pick for the eigenvectors, one of them I should have reversed the sign.

So, from this, I can't tell whether the eigenvector or its negative is the right one to use in there.

So the right way to do it is to, having settled on the signs, the Vs also, I don't know which sign to choose, but I choose one. I choose one.

And then, instead, I should have used the one that tells me what sign to choose, the rule that A times a V is σ times the U . So, having decided on the V , I multiply by A , I'll notice the factor σ coming out, and there will be a unit vector there, and I now know exactly what it is, and not only up to a change of sign. So that's the good and, of course, this is the main point about the SVD.

That's the point that we've diagonalized, that's A times the matrix of Vs equals U times the diagonal matrix of sigmas. That's the same as that.

OK. So that's, like, correcting the wrong sign from that earlier lecture.

And that would complete that, so that's how you would compute the SVD. Now, on the quiz, I going to ask -- well, maybe on the final.

So we've got quiz and final ahead.

Sometimes, you might be asked to find the SVD if I give you the matrix -- let me come back, now, to the main board -- or, I might give you the pieces. And I might ask you something about the matrix. For example, suppose I ask you, oh, let's say, if I tell you what σ is -- OK.

Let's take one example. Suppose σ is -- so all that's how we would compute them.

But now, suppose I give you these. Suppose I give you σ is, say, three two.

And I tell you that U has a couple of columns, and V has a couple of columns. OK.

Those are orthogonal columns, of course, because U and V are orthogonal. I'm just sort of, like, getting you to think about the SVD, because we only had that one lecture about it, and one homework, and, what kind of a matrix have I got here? What do I know about this matrix? All I really know right now is that its singular values, those sigmas are three and two, and the only thing interesting that I can see in that is that they're not zero. I know that this matrix is non-singular, right? That's invertible, I don't have any zero eigenvalues, and zero singular values, that's invertible, there's a typical SVD for a nice two-by-two non-singular invertible good matrix. If I actually gave you a matrix, then you'd have to find the U s and the V s as we just spoke. But, there.

Now, what if the two wasn't a two but it was -- well, let me make an extreme case, here -- suppose it was minus five. That's wrong, right away. That's not a singular value decomposition, right? The singular values are not negative.

So that's not a singular value decomposition, and forget it. OK.

So let me ask you about that one.

What can you tell me about that matrix? It's singular, right? It's got a singular matrix there in the middle, and, let's see, so, OK, it's singular, maybe you can tell me, its rank? What's the rank of A ? It's clearly -- somebody just say it -- one, thanks.

The rank is one, so the null space, what's the dimension of the null space? One. Right? We've got a two-by-two matrix of rank one, so of all that stuff from the beginning of the course is still with us.

The dimensions of those fundamental spaces is still central, and a basis for them. Now, can you tell me a vector that's in the null space? And then that will be my last point to make about the SVD. Can you tell me a vector that's in the null space? So what would I multiply by and get zero, here? I think the answer is probably v_2 . I think probably v_2 is in the null space, because I think that must be the eigenvector going with this zero eigenvalue. Yes.

Have a look at that. And I could ask you the null space of A transpose. And I could ask you the column space. All that stuff.

Everything is sitting there in the SVD. The SVD takes a little more time to compute, but it displays all the good stuff about a matrix.

OK. Any question about the SVD? Let me keep going with further topics.

Now, let's see. Similar matrices we've talked about, let me see if I've got another, -- OK. Here's a true false, so we can do that, easily. So.

Question, A given. A is symmetric and orthogonal.

OK. So beautiful matrices like that don't come along every day. But what can we say first about its eigenvalues? Actually, of course.

Here are our two most important classes of matrices, and we're looking at the intersection.

So those really are neat matrices, and what can you tell me about what could the possible eigenvalues be? Eigenvalues can be what? What do I know about the eigenvalues of a symmetric matrix? λ is real. What do I know about the eigenvalues of an orthogonal matrix? Ha. Maybe nothing.

But, no, that can't be. What do I know about the eigenvalues of an orthogonal matrix? Well, what feels right? Basing mathematics on just a little gut instinct here, the eigenvalues of an orthogonal matrix ought to have magnitude one.

Orthogonal matrices are like rotations, they're not changing the length, so orthogonal, the eigenvalues are one.

Let me just show you why. Why? So the matrix, can I call it Q for orthogonal for the moment? If I look at $Qx = \lambda x$, how do I see that this thing has magnitude one? I take the length of both sides.

This is taking lengths, taking lengths, this is whatever the magnitude is times the length of x .

And what's the length of Qx if Q is an orthogonal matrix? This is something you should know.

It's the same as the length of x .

Orthogonal matrices don't change lengths.

So λ has to be one. Right.

OK. That's worth committing to memory, that could show up again. OK. So what's the answer now to this question, what can the eigenvalues be? There's only two possibilities, and they are one and the other one, the other possibility is negative one, right, because these have the right magnitude, and they're real. OK.

TK. true -- OK.

True or false? A is sure to be positive definite. Well, this is a great matrix, but is it sure to be positive definite? No.

If it could have an eigenvalue minus one, it wouldn't be positive definite. True or false, it has no repeated eigenvalues. That's false, too. In fact, it's going to have repeated eigenvalues if it's as big as three by three, one of these $-$ one of these, at least, will have to get repeated.

Sure. So it's got repeated eigenvalues, but, is it diagonalizable? It's got these many, many, repeated eigenvalues.

If it's fifty by fifty, it's certainly got a lot of repetitions. Is it diagonalizable? Yes. All symmetric matrices, all orthogonal matrices can be diagonalized.

And, in fact, the eigenvectors can even be chosen orthogonal. So it could be, sort of, like, diagonalized the best way with a Q , and not just any old S . OK.

Is it non-singular? Is a symmetric orthogonal matrix non-singular? Sure.

Orthogonal matrices are always non-singular.

And, obviously, we don't have any zero Eigenvalues. Is it sure to be diagonalizable? Yes.

Now, here's a final step -- show that one-half of A plus I is A -- that is, prove one-half of A plus I is a projection matrix. OK? Let's see. What do I do? I could see two ways to do this.

I could check the properties of a projection matrix, which are what? A projection matrix is symmetric. Well, that's certainly symmetric, because A is. And what's the other property? I should square it, and hopefully get the same thing back. So can I do that, square and see if I get the same thing back? So if I square it, I'll get one-quarter of A squared plus two A plus I , right? And the question is, does that agree with p - the thing itself? One-half A plus I .

Hm. I guess I'd like to know something about A squared. What is A squared? That's our problem. What is A squared? If A is symmetric and orthogonal, A is symmetric and orthogonal. This is what we're given, right? It's symmetric, and it's orthogonal. So what's A squared? I . A squared is I , because A times A -- if A equals its own

inverse, so A times A is the same as A times A inverse, which is I . So this A squared here is I .

And now we've got it. We've got two identities over four, that's good, and we've got two A s over four, that's good. OK.

So it turned out to be a projection matrix safely.

And we could also have said, well, what are the eigenvalues of this thing? What are the eigenvalues of a half A plus I ? If the eigenvalues of A are one and minus one, what are the eigenvalues of A plus I ? Just stay with it these last thirty seconds here. What if I know these eigenvalues of A , and I add the identity, the eigenvalues of A plus I are zero and two.

And then when I divide by two, the eigenvalues are zero and one. So it's symmetric, it's got the right eigenvalues, it's a projection matrix.

OK, you're seeing a lot of stuff about eigenvalues, and special matrices, and that's what the quiz is about. OK, so good luck on the quiz.