DAVID SHIROKOFF: Hi everyone. Welcome back.

So today I'd like to tackle a problem on pseudoinverses. So given a matrix A, which is not square, so it's just 1 and 2. First, what is its pseudoinverse? So A plus I'm using to denote the pseudoinverse. Then secondly, compute A plus A and A A plus. And then thirdly, if x is in the null space of A, what is A plus A acting on x? And lastly, if x is in the column space of A transpose, what is A plus A*x?

So I'll let you think about this problem for a bit, and I'll be back in a second.

Hi everyone. Welcome back. OK, so let's take a look at this problem. Now first off, what is a pseudoinverse? Well, we define the pseudoinverse using the SVD. So in actuality, this is nothing new. Now, we note that because A is not square, the regular inverse of A doesn't necessarily exist. However, we do know that the SVD exists for every matrix A whether it's square or not.

So how do we compute the SVD of a matrix? Well let's just recall that the SVD of a matrix has the form of U sigma V transpose, where U and V are orthogonal matrices and sigma is a matrix with positive values along the diagonal or 0's along the diagonal. And let's just take a look at the dimensions of these matrices for a second. So we know that A is a 1 by 2 matrix.

And the way to figure out what the dimensions of these matrices are I usually always start with the center matrix, sigma, and sigma is always going to have the same dimensions as A, so it's going to be a 1 by 2 matrix. U and V are always square matrices. So to make this multiplication work out, we need V to have 2, and because it's square it has to be 2 by 2. And likewise, U has to be 1 by 1.

So we now have the dimensions of U, sigma, and V. And note, because U is a 1 by 1 matrix, the only orthogonal 1 by 1 matrix is just 1. So u we already know is just going to be the matrix, the identity matrix, which is a 1 by 1 matrix.

OK, now how do we compute V and sigma? Well, we can take A transpose and A, and if we do that, we end up getting the matrix V sigma transpose sigma V transpose. And this matrix is going to be a square matrix where the diagonal elements are squares of the singular values.

So computing V and the values along sigma just boil down to diagonalizing A transpose A.

So what is A transpose A? Well, in our case is [1; 2] times [1, 2], which gives us [1, 2; 2, 4]. And note that the second row is just a constant multiple times the first row.

Now what this means is we have a zero eigenvalue. So we already know that lambda_1 is going to be 0. So one of the eigenvalues of this matrix is 0. And of course, when we square root it, this is going to give us a singular value sigma, which is also 0. And this is generally a case when we have a sigma which is not square. We typically always have zero singular values.

Now to compute the second eigenvalue, well we already know how to compute the eigenvalues of a matrix, so I'm just going to tell you what it is. The second one is lambda is 5. And if we just take a quick look what the corresponding eigenvector is going to be to lambda is 5, it's going to satisfy this equation. So we can take the eigenvector u to be 1 and 2.

However, remember that when we compute the eigenvector for this orthogonal matrix V, they always have to have a unit length. And this vector right now doesn't have a unit length. We have to divide by the length of this vector, which in our case is 1 over root 5. And if I go back to the lambda equals 0 case, we also have another eigenvector, which I'll just state. You can actually compute it quite quickly just by noting that it has to be orthogonal to this eigenvector, 2 and 1.

So what this means is A has a singular value decomposition, which looks like: 1, so this is u, times sigma, which is going to be root 5, 0. Remember that the first sigma is actually the square root of the eigenvalue. Times a matrix which looks like, now we have to order the eigenvalues up in the correct order. Because 5 appears in the first column, we have to take this vector to be in the first column as well. So this is 1 over root 5, this is 2 over root 5, negative 2 over root 5, and 1 over root 5. And now this is V, but the singular value decomposition is defined by V transpose.

So this gives us a representation for A. And now once we have the SVD of A, how do we actually compute A plus, or the pseudoinverse of A? Well just note if A was invertible, then the inverse of A in terms of the SVD would be V transpose times the inverse of sigma. Sorry, this is not V transpose, this is just V. So it'd be V sigma inverse U transpose. And when A is invertible, sigma inverse exists.

So in our case, sigma inverse doesn't necessarily exist because sigma-- note, this is sigma-- sigma is root 5 and 0. So we have to construct a pseudoinverse for sigma. So the way that we do that is we take 1 over each singular value, and we take the transpose of sigma. So when A is not invertible, we can still construct a pseudoinverse by taking V, an approximation for sigma inverse, which in our case is going to be 1 over the singular value and 0. So note where sigma is invertible, we take the inverse, and then we fill in 0's in the other areas. Times U transpose.

And we can work this out. We get 1 over root 5, 1, minus 2; 2, 1, 1 over root 5, 0. And if I multiply things out, I get 1/5, [1; 2]. So this is an approximation for A inverse, which is the pseudoinverse.

So this finishes up part one. And I'll started on part two in a second.

So now that we've just computed A plus, the pseudoinverse of A. We're going to investigate some properties of the pseudoinverse. So for part two we need to compute A times A plus and A plus times A. So we can just go ahead and do this. So A A plus you can do fairly quickly. 1/5, [1; 2]. And when we multiply it out we get 1 plus 4 divided by 5 is 1. So we just get the one by one matrix, which is 1, the identity matrix.

And secondly, if we take A plus times A we're going to get 1/5, [1; 2] times [1, 2]. And we can just fill in this matrix. This is 1/5, [1, 2; 2, 1]. And this concludes part two.

So now let's take a look at what happens when a vector x is in the null space of A, and then secondly, what happens when x is in the column space of A transpose.

So for part three, let's assume x is in the null space of A. Well what's the null space of A? We can quickly check that the null space of A is a constant times any vector minus 2, 1.

So that's the null space. So if x is, for example, i.e. if we take x is equal to minus 2, 1, and we were to, say, multiply it by A plus A, acting on x, we see that we get 0. And this isn't very surprising because, well, if x is in the null space of A, we know that A acting on x is going to be 0. So that no matter what matrix A plus is, when we multiply by 0, we'll always end up with 0.

And then lastly, let's take a look at the column space of A transpose. Well, A transpose is [1, 2], so it's any constant times the vector [1; 2]. And specifically, if we were to take, say, x is equal to [1; 2], we can work at A plus A acting on the vector [1; 2]. So we have 1/5 [1, 2; 2, 1]. So recall this is A plus A. And if we multiply it on the vector [1; 2], we get 1 plus 4 is 5, divided by 5, so we get 1. 2 plus 2 is 4-- sorry, I copied the matrix down. So it's 2 plus 8, which is 10, divided by 5 is 2. And we see that at the end we recover the vector x.

So in general, if we take A plus A acting on x, where x is in the column space of A transpose, we always recover x at the end of the day. So intuitively, what does this matrix A plus A do? Well, if x is in the null space of A, it just kills it. We just get 0. If x is not in the null space of A, then we just get x back. So it's essentially the identity matrix acting on x whenever x is in the column space of A transpose.

Now specifically, if A is invertible, then A doesn't have a null space. So what that means is: when A is invertible, A plus A recovers the identity because when we multiply it on any vector, we get that vector back.

So I'd like to conclude here, and I'll see you next time.