Eigenvalues, Eigenvectors, Diagonalization and Choice of Basis

Justin Werfel

9.29 Optional Lecture #3, 2/24/03

1 Fundamentals

What do we mean when we write a vector in a form like $x = [2, -3]^T$? What do these components mean? Well, we've got a two-dimensional space; and we choose basis vectors that span it: that is, any vector in this space can be written as a linear combination of those basis vectors. Typically we choose the basis vectors $\hat{e}_1$ and $\hat{e}_2$, of unit length[1] in, by convention, the x- and y-directions. So writing $x = [2, -3]^T$ is shorthand for $x = 2\hat{e}_1 - 3\hat{e}_2$. And we could choose different basis vectors; for instance, if we choose $e_3 = 1\hat{e}_1 + \frac{1}{2}\hat{e}_2$ and $e_4 = \frac{1}{2}\hat{e}_1 + 1\hat{e}_2$, then $x = \frac{14}{3}e_3 - \frac{16}{3}e_4$. The basis vectors don't have to be orthogonal or have unit length, either; they just have to span the 2-D space. We almost always use $\hat{e}_1$ and $\hat{e}_2$ because they're the most convenient, but there's nothing intrinsically more fundamental about them.

Now let's say you've got a square matrix

$$A = \begin{bmatrix} 5/3 & 2/3 \\ -2/3 & 10/3 \end{bmatrix}$$

What happens to a vector when you multiply it by A? Let's use one of our basis vectors, say $\hat{e}_1$. If you do the multiplication, you find that $A\hat{e}_1$ is $[5/3, -2/3]^T$. The original vector got both rotated and scaled; and you need both basis vectors to describe the result.

Now suppose we try it with one of these alternate basis vectors. What's Ae_3? It turns out to be $[2, 1]^T$—which is $2e_3$. Applying A to this vector resulted in it being scaled, but not rotated. That's a special case. While we're at it, let's see what happens to e_4: $Ae_4 = [3/2, 3]^T = 3e_4$. Same thing: the result is just a multiple of the original vector.

These vectors that, when you multiply a matrix by them, give you back a multiple of themselves, are called *eigenvectors* of that matrix. The multiple they give is the corresponding *eigenvalue*.[2]

Why do we care about these things? Let me just give you an idea of that now, and then I'll come back to it a little bit later. Take the vector x we started with, that we

---

[1] The hat over the vector name signifies that it has unit length.

[2] 'Eigen' is German for something like 'own'.

can write with reference to our original basis $\hat{e}_1$ and $\hat{e}_2$ as $2\hat{e}_1 - 3\hat{e}_2$. We apply A to it and get $Ax = 2A\hat{e}_1 - 3A\hat{e}_2 = 2(\frac{5}{3}\hat{e}_1 - \frac{2}{3}\hat{e}_2) - 3(\frac{2}{3}\hat{e}_1 - \frac{10}{3}\hat{e}_2) = \frac{4}{3}\hat{e}_1 - 34/3\hat{e}_2$. For comparison, suppose we express x with reference to the basis of vectors we've just said are eigenvectors: $x = \frac{14}{3}e_3 - \frac{16}{3}e_4$. Then $Ax = \frac{14}{3}Ae_3 - \frac{16}{3}Ae_4 = \frac{14}{3}\cdot 2e_3 - \frac{16}{3}\cdot 3e_4 = \frac{28}{3}e_3 - 16e_4$. Each component is just scaled separately by its corresponding eigenvalue. So if we express things in the eigenvector basis, things are much neater and our life is easier.

2 Finding the eigenvalues and eigenvectors

The basic eigenvalue equation is $Ax = \lambda x$; you apply A to the eigenvector x, and get back the eigenvector scaled by the eigenvalue λ. So $(A - I\lambda)x = 0$.[3] That means that x is in the nullspace of $(A - I\lambda)$; there's a nonzero vector that, when multiplied by $(A - I\lambda)$, gets mapped to zero. That means that $(A - I\lambda)$ must be singular, and its determinant must be zero. So we take the determinant of $(A - I\lambda)$, set it to zero, and solve for λ. $|A - I\lambda| = 0$ is called the *characteristic equation*. And since $(A - I\lambda)$ is $n \times n$, there are n roots, and that gives you the n eigenvalues.

In this example,

$$\begin{aligned} A - I\lambda &= \begin{bmatrix} 5/3 - \lambda & 2/3 \\ -2/3 & 10/3 - \lambda \end{bmatrix}; \\ |A - I\lambda| &= 50/9 - 15/3 * \lambda + \lambda^2 + 4/9 = 0 \\ \Rightarrow \lambda = (5 \pm sqrt(225/9 - 4 * 54/9))/2 &= (5 \pm sqrt(1))/2 = \{2, 3\}. \end{aligned}$$

To find the eigenvectors, you substitute each of these eigenvalues back into the original equation, and solve the resulting system of equations. First, $\lambda = 2$:

$$\begin{gathered} (A - I2)x = 0 \Rightarrow \\ \begin{bmatrix} -1/3 & 2/3 \\ -2/3 & 4/3 \end{bmatrix} x = 0; \end{gathered}$$

if you write $x = [x_1 \quad x_2]^T$, that means $-x_1/3 + 2x_2/3 = 0 \Rightarrow x_1 = 2x_2$, for instance, $x = [1, .5]^T$. (That's our vector e_3, corresponding to the eigenvalue 2; so that checks out.) For $\lambda = 3$,

$$\begin{gathered} (A - I3)x = 0 \Rightarrow \\ \begin{bmatrix} -4/3 & 2/3 \\ -2/3 & 1/3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 0 \Rightarrow 2x_1 = x_2, \end{gathered}$$

e.g., $x = [.5, 1]^T = e_4$.

For each of these eigenvalues, there's no single corresponding eigenvector. Any vector in this direction (or in the opposite direction) works equally well. And you

---

[3] Note that we're not just writing $(A - \lambda)x = 0$. This is a matter of notation, so that all our matrix arithmetic is well-defined; subtracting a scalar from a matrix is not well-defined in general. *In particular*, note that if you write `A-lambda` in MATLAB, `lambda` will be subtracted from *every* entry of `A`, whereas we want it to be subtracted only from the entries on the main diagonal.

can see that by $A(kx) = k(Ax) = k(\lambda x) = \lambda(kx)$: so if x is an eigenvector with eigenvalue λ, so is kx, for nonzero k.

Here are a few useful facts about eigenvalues that I'm going to tell you without proof.

- The product of the eigenvalues of A is the determinant of A: $\prod_i \lambda_i = |A|$.
- The sum of the eigenvalues is the trace of the matrix, i.e. the sum of the diagonal elements: $\sum_i \lambda_i = \text{tr}(A) = \sum_i A_{ii}$.
- If all the eigenvalues are different, the eigenvectors are linearly independent (though the converse is not necessarily true).

3 Diagonalization

Now that we have the eigenvectors, what do we do with them? We can use them to *diagonalize* the original matrix, as it's called. If we put all the eigenvectors together into a matrix S whose columns are the eigenvectors $\{x_1, x_2, \ldots\}$[4], then we'll have $AS = A[x_1, x_2, \ldots] = [\lambda_1 x_1, \lambda_2 x_2, \ldots] = S * \text{diag}(\lambda_1, \lambda_2, \ldots)$. That last matrix is called Λ (capital λ); it's the diagonal matrix with the eigenvalues on the main diagonal and 0s everywhere else. So we have $AS = S\Lambda$. If the eigenvectors are all linearly independent, then S will be invertible, and we can write $S^{-1}AS = \Lambda$. Or equivalently, $S\Lambda S^{-1} = A$. This is diagonalization.

Why would we want to diagonalize a matrix? Well, for one thing, it makes many operations trivial. It's a real pain to raise an arbitrary matrix to a power or to find its inverse, but with a diagonal matrix, it's one step. (And so here, $A^{-1} = (S\Lambda S^{-1})^{-1} = S\Lambda^{-1}S^{-1}$; and $A^2 = (S\Lambda S^{-1})(S\Lambda S^{-1}) = S\Lambda^2 S^{-1}$ (and by extension to higher powers).)

Diagonalization is one of many useful ways to factorize a matrix. Here we've factorized it into this useful combination of meaningful and simple matrices. Or you can think of diagonalization as a way of choosing a useful basis for the vector space, such that if A is written in that basis, it's diagonal. This technique of choosing the best basis can be extended to arbitrary matrices, not necessarily square, in singular value decomposition. I think we're going to get to that later in the course, but the idea is to choose bases that are orthonormal that also diagonalize the matrix. The standard basis $\{\hat{e}_1, \hat{e}_2\}$ is orthonormal, that is, all the vectors are of length 1 and perpendicular to all the others, but in that basis the matrix typically isn't diagonal. The eigenvalue basis diagonalizes the matrix, but as in the example we've been using, the vectors aren't all orthogonal. So in singular value decomposition, you end up having to use two separate sets of bases: $A = U\Sigma V^T$. (U and V are both orthogonal, that is, $U^T = U^{-1}$ and $V^T = V^{-1}$.) But I'm not going to get into that now.

To confuse things further, there's a distinction between *right* and *left* eigenvectors. If you multiply A on the right by a column vector x and get back a multiple of x, $Ax = \lambda x$, that's a right eigenvector. (And that's the kind you always hear about, and what people are referring to if they just say 'eigenvector'.) But also, if you multiply A on

---

[4] The subscripts here index the different eigenvectors, not entries of a single eigenvector.

the left by a row vector y^T and get back a multiple of $y^T, y^T A = \lambda y^T$, that's a left eigenvector. And starting from the equation $A = S\Lambda S^{-1}$, you can multiply on the right by S and get $AS = S\Lambda$, so S is the matrix whose columns are the right eigenvectors; or you can multiply on the left by S^{-1} and get $S^{-1}A = \Lambda S^{-1}$, so S^{-1} is the matrix whose rows are the left eigenvectors.

In general the left and right eigenvectors can be different (though their eigenvalues will be the same—Λ is the same in both equations above), and the eigenvectors need not be orthogonal. But if you have $AA^T = A^T A$—this is an important special case—then the eigenvectors of A will be perpendicular, and since you can choose the magnitude of the eigenvectors arbitrarily, the matrix S can be chosen to be orthonormal. In that case, we call it Q to emphasize this special case: $A = Q\Lambda Q^{-1}$ = (since Q is orthogonal, $Q^T = Q^{-1}$) $Q\Lambda Q^T$. Here you can see the left and right eigenvectors are the same: $AQ = Q\Lambda$, $Q^T A = \Lambda Q^T$.

This class of matrices for which $AA^T = A^T A$ includes:

- *symmetric* matrices, $A^T = A$
  - If A is real and symmetric, its eigenvalues will all be real.
- *skew-symmetric* matrices, $A^T = -A$
  - If A is real and skew-symmetric, its eigenvalues will be purely imaginary.
- *orthogonal* matries, $A^T = A^{-1}$
  - All the eigenvalues of an orthogonal matrix have magnitude 1.

4 The Eigenvector Game

Can you find the educational value hidden in `nnd6eg`? I can't.