A useful basis for defective matrices: Generalized eigenvectors and the Jordan form

S. G. Johnson, MIT 18.06 Spring 2009

April 22, 2009

1 Introduction

So far in the eigenproblem portion of 18.06, our strategy has been simple: find the eigenvalues λ_i and the corresponding eigenvectors $\mathbf{x}_i$ of a square matrix A, expand any vector of interest $\mathbf{u}$ in the basis of these eigenvectors ($\mathbf{u} = c_1\mathbf{x}_1 + \cdots + c_n\mathbf{x}_n$), and then any operation with A can be turned into the corresponding operation with λ_i acting on each eigenvector. So, A^k becomes λ_i^k, e^{At} becomes $e^{\lambda_i t}$, and so on. But this relied on one key assumption: we require the $n \times n$ matrix A to have a *basis* of n independent eigenvectors. We call such a matrix A **diagonalizable**.

Many important cases are always diagonalizable: matrices with n distinct eigenvalues λ_i, real symmetric or orthogonal matrices, and complex Hermitian or unitary matrices. But there are rare cases where A does *not* have a complete basis of n eigenvectors: such matrices are called **defective**. For example, consider the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

This matrix has a characteristic polynomial $\lambda^2 - 2\lambda + 1$, with a repeated root (a single eigenvalue) $\lambda_1 = 1$. (Equivalently, since A is upper triangular, we can read the determinant of $A - \lambda I$, and hence the eigenvalues, off the diagonal.) However, it only has a *single* indepenent eigenvector, because

$$A - I = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

is obviously rank 1, and has a one-dimensional nullspace spanned by $\mathbf{x}_1 = (1, 0)$.

Defective matrices arise rarely in practice, and usually only when one arranges for them intentionally, so we have not worried about them up to now. But it is important to have some idea of what happens when you have a defective matrix. Partially for computational purposes, but also to understand conceptually what is possible. For example, what will be the result of

$$A^k \begin{pmatrix} 1 \\ 2 \end{pmatrix} \quad \text{or} \quad e^{At} \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

for the defective matrix A above, since $(1, 2)$ is not in the span of the (single) eigenvector of A? For diagonalizable matrices, this would grow as λ^k or $e^{\lambda t}$, respectively, but what about defective matrices?

The textbook (*Intro. to Linear Algebra, 4th ed.* by Strang) covers the defective case only briefly, in section 6.6, with something called the **Jordan form** of the matrix, a generalization of diagonalization. In that short section, however, the Jordan form falls down out of the sky, and you don't really learn where it comes from or what its consequences are. In this section, we will take a different tack. For a diagonalizable matrix, the fundamental vectors are the eigenvectors, which are useful in their own right and give the diagonalization of the matrix as a side-effect. For a defective matrix, to get a complete basis we need to supplement the eigenvectors with something called **generalized eigenvectors**. Generalized eigenvectors are useful in their own right, just like eigenvectors, and also give the Jordan form as a side effect. Here, however, we'll focus on defining, obtaining, and using the generalized eigenvectors, and not so much on the Jordan form.

2 *Defining* generalized eigenvectors

In the example above, we had a 2×2 matrix A but only a single eigenvector $\mathbf{x}_1 = (1,0)$. We need another vector to get a basis for $\mathbb{R}^2$. Of course, we could pick another vector at random, as long as it is independent of $\mathbf{x}_1$, but we'd like it to have something to do with A, in order to help us with computations just like eigenvectors. The key thing is to look at $A-I$ above, and to notice that the one of the columns is equal to $\mathbf{x}_1$: the vector $\mathbf{x}_1$ is in $C(A-I)$, so we can find a new **generalized eigenvector** $\mathbf{x}_1^{(2)}$ satisfying

$$(A-I)\mathbf{x}_1^{(2)} = \mathbf{x}_1, \qquad \mathbf{x}_1^{(2)} \perp \mathbf{x}_1.$$

Notice that, since $\mathbf{x}_1 \in N(A-I)$, we can add any multiple of $\mathbf{x}_1$ to $\mathbf{x}_1^{(2)}$ and still have a solution, so we can use Gram-Schmidt to get a *unique* solution $\mathbf{x}_1^{(2)}$ perpendicular to $\mathbf{x}_1$. This particular 2×2 equation is easy enough for us to solve by inspection, obtaining $\mathbf{x}_1^{(2)} = (0,1)$. Now we have a nice *orthonormal* basis for $\mathbb{R}^2$, and our basis has some simple relationship to A!

Before we talk about how to *use* these generalized eigenvectors, let's give a more general definition. Suppose that λ_i is an eigenvalue of A corresponding to a repeated root of $\det(A-\lambda_i I)$, but with only a single (ordinary) eigenvector $\mathbf{x}_i$, satisfying, as usual:

$$(A-\lambda_i I)\mathbf{x}_i = 0.$$

If λ_i is a double root, we will need a second vector to complete our basis. Remarkably,[1] it turns out to *always* be the case for a double root λ_i that $\mathbf{x}_i$ is in $C(A-\lambda_i I)$, just as in the example above. Hence, we can *always* find a unique second solution $\mathbf{x}_i^{(2)}$ satisfying:

$$\boxed{(A-\lambda_i I)\mathbf{x}_i^{(2)} = \mathbf{x}_i, \qquad \mathbf{x}_i^{(2)} \perp \mathbf{x}_i}.$$

Again, we can choose $\mathbf{x}_i^{(2)}$ to be perpendicular to $\mathbf{x}_i^{(1)}$ via Gram-Schmidt—this is not strictly necessary, but gives a convenient orthogonal basis. (That is, the complete solution is always of the form $\mathbf{x}_p + c\mathbf{x}_i$, a particular solution $\mathbf{x}_p$ plus any multiple of the nullspace basis $\mathbf{x}_i$. If we choose $c = -\mathbf{x}_i^T\mathbf{x}_p/\mathbf{x}_i^T\mathbf{x}_i$ we get the unique orthogonal solution $\mathbf{x}_i^{(2)}$.) We call $\mathbf{x}_i^{(2)}$ a **generalized eigenvector** of A.

[1] This fact is proved in any number of advanced textbooks on linear algebra; I like *Linear Algebra* by P. D. Lax.

2.1 More than double roots

If we wanted to be more notationally consistent, we could use $\mathbf{x}_i^{(1)}$ instead of $\mathbf{x}_i$. If λ_i is a triple root, we would find a third vector $\mathbf{x}_i^{(3)}$ perpendicular to $\mathbf{x}_i^{(1,2)}$ by requiring $(A-\lambda_i I)\mathbf{x}_i^{(3)} = \mathbf{x}_i^{(2)}$, and so on. In general, if λ_i is an m-times repeated root, then we will always be able to find an orthogonal sequence of generalized eigenvectors $\mathbf{x}_i^{(j)}$ for $j = 2\ldots m$ satisfying $(A-\lambda_i I)\mathbf{x}_i^{(j)} = \mathbf{x}_i^{(j-1)}$ and $(A-\lambda_i I)\mathbf{x}_i^{(1)} = 0$. Even more generally, you might have cases with e.g. a triple root and two ordinary eigenvectors, where you need only one generalized eigenvector, or an m-times repeated root with $\ell > 1$ eigenvectors and $m-\ell$ generalized eigenvectors. However, cases with more than a double root are extremely rare in practice. Defective matrices are rare enough to begin with, so here we'll stick with the most common defective matrix, one with a double root λ_i: hence, one ordinary eigenvector $\mathbf{x}_i$ and one generalized eigenvector $\mathbf{x}_i^{(2)}$.

3 *Using* generalized eigenvectors

Using an eigenvector $\mathbf{x}_i$ is easy: multiplying by A is just like multiplying by λ_i. But how do we use a generalized eigenvector $\mathbf{x}_i^{(2)}$? The key is in the definition above. It immediately tells us that

$$A\mathbf{x}_i^{(2)} = \lambda_i \mathbf{x}_i^{(2)} + \mathbf{x}_i.$$

It will turn out that this has a simple consequence for more complicated expressions like A^k or e^{At}, but that's probably not obvious yet. Let's try multiplying by A^2:

$$\begin{aligned} A^2\mathbf{x}_i^{(2)} &= A(A\mathbf{x}_i^{(2)}) = A(\lambda_i\mathbf{x}_i^{(2)} + \mathbf{x}_i) = \lambda_i(\lambda_i\mathbf{x}_i^{(2)} + \mathbf{x}_i) + \lambda_i\mathbf{x}_i \\ &= \lambda_i^2\mathbf{x}_i^{(2)} + 2\lambda_i\mathbf{x}_i \end{aligned}$$

and then try A^3:

$$\begin{aligned} A^3\mathbf{x}_i^{(2)} &= A(A^2\mathbf{x}_i^{(2)}) = A(\lambda_i^2\mathbf{x}_i^{(2)} + 2\lambda_i\mathbf{x}_i) = \lambda_i^2(\lambda_i\mathbf{x}_i^{(2)} + \mathbf{x}_i) + 2\lambda_i^2\mathbf{x}_i \\ &= \lambda_i^3\mathbf{x}_i^{(2)} + 3\lambda_i^2\mathbf{x}_i. \end{aligned}$$

From this, it's not hard to see the general pattern (which can be formally proved by induction):

$$\boxed{A^k\mathbf{x}_i^{(2)} = \lambda_i^k\mathbf{x}_i^{(2)} + k\lambda_i^{k-1}\mathbf{x}_i}.$$

Notice that the coefficient in the second term is exactly $\frac{d}{d\lambda_i}(\lambda_i)^k$. This is the clue we need to get the general formula to apply any function $f(A)$ of the matrix A to the eigenvector and the generalized eigenvector:

$$f(A)\mathbf{x}_i = f(\lambda_i)\mathbf{x}_i,$$

$$\boxed{f(A)\mathbf{x}_i^{(2)} = f(\lambda_i)\mathbf{x}_i^{(2)} + f'(\lambda_i)\mathbf{x}_i}.$$

Multiplying by a function of the matrix multiplies $\mathbf{x}_i^{(2)}$ by the same function of the eigenvalue, just as for an eigenvector, but *also* adds a term multiplying $\mathbf{x}_i$ by the *derivative* $f'(\lambda_i)$. So, for example:

$$\boxed{e^{At}\mathbf{x}_i^{(2)} = e^{\lambda_i t}\mathbf{x}_i^{(2)} + te^{\lambda_i t}\mathbf{x}_i}.$$

We can show this explicitly by considering what happens when we apply our formula for A^k in the Taylor series for e^{At}:

$$e^{At}\mathbf{x}_i^{(2)} = \sum_{k=0}^{\infty}\frac{A^k t^k}{k!}\mathbf{x}_i^{(2)} = \sum_{k=0}^{\infty}\frac{t^k}{k!}(\lambda_i^k\mathbf{x}_i^{(2)} + k\lambda_i^{k-1}\mathbf{x}_i)$$
$$= \sum_{k=0}^{\infty}\frac{(\lambda_i t)^k}{k!}\mathbf{x}_i^{(2)} + t\sum_{k=1}^{\infty}\frac{(\lambda_i t)^{k-1}}{(k-1)!}\mathbf{x}_i = e^{\lambda_i t}\mathbf{x}_i^{(2)} + te^{\lambda_i t}\mathbf{x}_i.$$

In general, that's how we show the formula for $f(A)$ above: we Taylor expand each term, and the A^k formula means that each term in the Taylor expansion has corresponding term multiplying $\mathbf{x}_i^{(2)}$ and a *derivative* term multiplying $\mathbf{x}_i$.

3.1 More than double roots

In the rare case of two generalized eigenvectors from a triple root, you will have a generalized eigenvector $\mathbf{x}_i^{(3)}$ and get a $f(A)\mathbf{x}_i^{(3)} = f(\lambda)\mathbf{x}_i^{(3)} + f'(\lambda)\mathbf{x}_i^{(2)} + f''(\lambda)\mathbf{x}_i$, where the f'' term will give you $k(k-1)\lambda_i^{k-2}$ and $t^2 e^{\lambda_i t}$ for A^k and e^{At} respectively. A quadruple root with one eigenvector and three generalized eigenvectors will give you f''' terms (that is, k^3 and t^3 terms), and so on. The theory is quite pretty, but doesn't arise often in practice so I will skip it; it is straightforward to work it out on your own if you are interested.

3.2 Example

Let's try this for our example 2×2 matrix $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ from above, which has an eigenvector $\mathbf{x}_1 = (1,0)$ and a generalized eigenvector $\mathbf{x}_1^{(2)} = (0,1)$ for an eigenvalue $\lambda_1 = 1$. Suppose we want to comput $A^k\mathbf{u}_0$ and $e^{At}\mathbf{u}_0$ for $\mathbf{u}_0 = (1,2)$. As usual, our first step is to write $\mathbf{u}_0$ in the basis of the eigenvectors...except that now we also include the generalized eigenvectors to get a complete basis:

$$\mathbf{u}_0 = \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \mathbf{x}_1 + 2\mathbf{x}_1^{(2)}.$$

Now, computing $A^k\mathbf{u}_0$ is easy, from our formula above:

$$\begin{aligned} A^k\mathbf{u}_0 &= A^k\mathbf{x}_1 + 2A^k\mathbf{x}_1^{(2)} = \lambda_1^k\mathbf{x}_1 + 2\lambda_1^k\mathbf{x}_1^{(2)} + 2k\lambda_1^{k-1}\mathbf{x}_1 \\ &= 1^k\begin{pmatrix} 1 \\ 2 \end{pmatrix} + 2k\,1^{k-1}\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1+2k \\ 2 \end{pmatrix}. \end{aligned}$$

For example, this is the solution to the recurrence $\mathbf{u}_{k+1} = A\mathbf{u}_k$. Even though A has only an eigenvalue $|\lambda_1| = 1 \le 1$, the solution still blows up, but it blows up *linearly* with k instead of exponentially.

Consider instead $e^{At}\mathbf{u}_0$, which is the solution to the system of ODEs $\frac{d\mathbf{u}(t)}{dt} = A\mathbf{u}(t)$ with the initial condition $\mathbf{u}(0) = \mathbf{u}_0$. In this case, we get:

$$\begin{aligned} e^{At}\mathbf{u}_0 &= e^{At}\mathbf{x}_1 + 2e^{At}\mathbf{x}_1^{(2)} = e^{\lambda_1 t}\mathbf{x}_1 + 2e^{\lambda_1 t}\mathbf{x}_1^{(2)} + 2te^{\lambda_1 t}\mathbf{x}_1 \\ &= e^t\begin{pmatrix} 1 \\ 2 \end{pmatrix} + 2te^t\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1+2t \\ 2 \end{pmatrix}e^t. \end{aligned}$$

In this case, the solution blows up exponentially since $\lambda_1 = 1 > 0$, but we have an *additional* term that blows up as an exponential multiplied by t.

Those of you who have taken 18.03 are probably familiar with these terms multiplied by t in the case of a repeated root. In 18.03, it is presented simply as a guess for the solution that turns out to work, but here we see that it is part of a more general pattern of generalized eigenvectors for defective matrices.

4 The Jordan form

For a diagonalizable matrix A, we made a matrix S out of the eigenvectors, and saw that multiplying by A was equivalent to multiplying by $S\Lambda S^{-1}$ where $\Lambda = S^{-1}AS$ is the diagonal matrix of eigenvalues, the *diagonalization* of A. Equivalently, $AS = \Lambda S$: A multiplies each column of S by the corresponding eigenvalue. Now, we will do exactly the same steps for a defective matrix A, using the basis of eigenvectors and generalized eigenvectors, and obtain the **Jordan form** J instead of Λ.

Let's consider a simple case of a 4×4 first, in which there is only *one* repeated root λ_2 with an eigenvector $\mathbf{x}_2$ and a generalized eigenvector $\mathbf{x}_2^{(2)}$, and the other two eigenvalues λ_1 and λ_3 are distinct with independent eigenvectors $\mathbf{x}_1$ and $\mathbf{x}_3$. Form a matrix $M = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_2^{(2)}, \mathbf{x}_3)$ from this basis of four vectors (3 eigenvectors and 1 generalized eigenvector). Now, consider what happends when we multiply A by M:

$$\begin{aligned} AM &= (\lambda_1 \mathbf{x}_1, \lambda_2 \mathbf{x}_2, \lambda_2 \mathbf{x}_2^{(2)} + \mathbf{x}_2, \lambda_3 \mathbf{x}_3). \\ &= M \begin{pmatrix} \lambda_1 & & & \\ & \lambda_2 & 1 & \\ & & \lambda_2 & \\ & & & \lambda_3 \end{pmatrix} = MJ. \end{aligned}$$

That is, $A = MJM^{-1}$ where J is *almost* diagonal: it has $\lambda's$ along the diagonal, but it *also has 1's above the diagonal for the columns corresponding to generalized eigenvectors*. This is exactly the Jordan form of the matrix A. J, of course, has the same eigenvalues as A since A and J are similar, but J is much simpler than A.

The generalization of this, when you perhaps have more than one repeated root, and perhaps the multiplicity of the root is greater than 2, is fairly obvious, and leads immediately to the formula given without proof in section 6.6 of the textbook. What I want to emphasize here, however, is not so much the formal theorem that a Jordan form exists, but how to *use* it via the generalized eigenvectors: in particular, that generalized eigenvectors give us *linearly growing* terms $k\lambda^{k-1}$ and $te^{\lambda t}$ when we multiply by A^k and e^{At}, respectively.