14.102, Math for Economists
Fall 2005
Lecture Notes, 9/27/2005

These notes are primarily based on those written by George Marios Angeletos for the Harvard Math Camp in 1999 and 2000, and updated by Stavros Panageas for the MIT Math for Economists Course in 2002. I have made only minor changes to the order of presentation, and added some material from Guido Kuersteiner's notes on linear algebra for 14.381 in 2002. The usual disclaimer applies; questions and comments are welcome.

Nathan Barczi
nab@mit.edu

6 Matrix Diagonalization and Eigensystems

6.1 The Characteristic Equation, Eigenvalues and Eigenvectors

In this section we study eigenvalues and eigenvectors of a given matrix A. These can be used to transform the matrix A into a simpler form which is useful for solving systems of linear equations and analyzing the properties of the mapping described by A. We say that λ is an **eigenvalue** of an $n \times n$ matrix A with corresponding **eigenvector** v if

$$Av = \lambda v$$

for some $v \neq 0$. Conversely, we say that $v \neq 0$ satisfying the equation is an eigenvector corresponding to the eigenvalue λ.

Note that we can rewrite the above equation as $(A - \lambda I)v = 0$. This homogenous system of equations has a nontrivial solution if and only if the matrix $A - \lambda I$ is singular, which in turn holds iff $|A - \lambda I| = 0$. This leads to a characterization of the eigenvalues as solutions to the equation $|A - \lambda I| = 0$. Note that $|A - \lambda I|$ is a polynomial of degree n in λ (why?). It thus has at most n, possibly complex, roots, and at least one.

For square matrices we define

Definition 153 *Let A be any $n \times n$ matrix and I the $n \times n$ identity matrix. The* ***characteristic polynomial*** *of A is $\xi(\lambda) \equiv |A - \lambda I|$. Its characteristic equation is $\xi(\lambda) = 0$, and the solutions to it are called* ***characteristic roots****, or* ***eigenvalues****. Any vector $v \neq 0$ that satisfies $Av = \lambda v$ for some λ is a* ***characteristic vector****, or* ***eigenvector****.*

Exercise 154 *Let a 2×2 matrix*

$$A = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$

Show that then $\xi(\lambda) \equiv |A - \lambda I| = \lambda^2 - (\alpha+\delta)\lambda + (\alpha\delta - \beta\gamma)$*, and find the eigenvalues.*

Exercise 155 *By using the inductive definition of the determinant show, or at least persuade yourself, that if* A *is* $n \times n$ *then* $\xi(\lambda) \equiv |A - \lambda I|$ *is an* n*-th order polynomial.*

By the fundamental theorem of algebra, any n-th order polynomial has exactly n roots. Of course, some of these roots may be imaginary rather than real, or they might be repeated. In any case, there are (complex or real) numbers $\{\lambda_1, ..., \lambda_n\}$ such that

$$\xi(\lambda) \equiv |A - \lambda I| = (\lambda - \lambda_1)....(\lambda - \lambda_n)$$

These $\{\lambda_1, ..., \lambda_n\}$ are the eigenvalues of A.

Remark: We emphasize that, in what follows, we do not assume the eigenvalues or the eigenvectors to be **real**. A real matrix A may well have nonreal eigenvalues and nonreal eigenvectors.

Nor do we assume that all roots are distinct. If a root appears once, then it is called a **distinct** root, while if it is repeated $r > 1$ times, then it is called a r**-fold** root. For example, the only eigenvalue of the identity matrix is 1, appearing with multiplicity n; i.e., it is an n-fold root. Moreover, in this example the eigenvectors are not unique, either. Indeed, all nonzero vectors $v \in \mathbb{R}^n$ are eigenvectors of the identity matrix associated to eigenvalue 1.

Returning to the eigenvectors of A, we observe that the singularity [since $\xi(\lambda_j) \equiv |A - \lambda_j I| = 0$] of matrix $(A - \lambda_j I)$, implies that there exists a vector $v_j \neq 0$ such that $(A - \lambda_j I)v_j = 0$. Rearranging we get $Av_j = \lambda_j v_j$. Thus:

Lemma 156 *To any eigenvalue of* A*,* λ_j *such that.* $|A - \lambda_j I| = 0$*, there is associated at least one eigenvector* $v_j \neq 0$ *such that* $Av_j = \lambda_j v_j$*.*

We showed that each eigenvalue has at least one eigenvector associated with it. In fact, it has a whole continuum of associated eigenvectors: Indeed, take $v_j \neq 0$ such that $Av_j = \lambda_j v_j$ and let $w_j = \mu v_j$ for any scalar $\mu \neq 0$. Then $w_j \neq 0$ and

$$Av_j = \lambda_j v_j \Rightarrow A(\mu v_j) = \lambda_j(\mu v_j) \Rightarrow Aw_j = \lambda_j w_j$$

meaning that w_j is as well an eigenvector associated with eigenvalue λ_j. Given this intrinsic multiplicity, we define

Definition 157 *Let an* $n \times n$ *matrix* A *and let* $\{\lambda_1, ..., \lambda_n\}$ *its eigenvalues. For each* λ_j *we define the corresponding* ***characteristic manifold*** *or* ***eigenspace*** $\mathbb{M}_j$ *as the subspace of all eigenvectors associated with* λ_j*:*[7]

$$\mathbb{M}_j \equiv \{x \in \mathbb{R}^n | Ax = \lambda_j x\}$$

[7] Plus the zero vector.

Exercise 158 *Show that*

$$v, w \in \mathbb{M}_j \Rightarrow \alpha v + \beta w \in \mathbb{M}_j \ \forall \alpha, \beta \in \mathbb{R}$$

which means that $\mathbb{M}_j$ *is indeed a vector (sub)space. Show also that it has* $\dim(\mathbb{M}_j) \geq 1$*, and notice that* $\mathbb{M}_j$ *is unique for each* λ_j*.*

A more interesting observation is the following:

Theorem 159 *Suppose* k $(k \leq n)$ *eigenvalues* $\{\lambda_1, ..., \lambda_k\}$ *of* A *are distinct, and take any corresponding eigenvectors* $\{v_1, ..., v_k\}$*, defined by* $v_j \neq 0, Av_j = \lambda_j v_j$ *for* $j = 1, ..., k$*. Then,* $\{v_1, ..., v_k\}$ *are linearly independent.*

Proof. *First consider two such eigenvectors. Suppose we have eigenvalue* λ *with eigenvector* v*, and eigenvalue* μ *with eigenvector* w*,* $\lambda \neq \mu$*. We will show that* $\alpha v + \beta w = 0 \Rightarrow \alpha = \beta = 0$*, implying that* v *and* w *are linearly independent.*

So suppose we have

$$\begin{aligned} \alpha v + \beta w &= 0 && (1) \\ \alpha A v + \beta A w &= 0 && (2) \\ \alpha \lambda v + \beta \mu w &= 0 && (3) \end{aligned}$$

Now multiply (1) by $-\lambda$ *and add to (3) to get*

$$\beta(\mu - \lambda)w = 0 \qquad (4)$$

which implies that $\beta = 0$*, and plugging this back into (1) implies that* $\alpha = 0$ *as well.*

Now consider any three eigenvectors $(v, w$ *and* $u)$ *with distinct eigenvalues* $(\lambda, \mu$ *and* $\nu)$*; we proceed in much the same manner:*

$$\begin{aligned} \alpha v + \beta w + \gamma u &= 0 && (5) \\ \alpha A v + \beta A w + \gamma A u &= 0 && (6) \\ \alpha \lambda v + \beta \mu w + \gamma \nu u &= 0 && (7) \\ \beta(\mu - \lambda)w + \gamma(\nu - \lambda)u &= 0 && (8) \end{aligned}$$

But this is a linear combination of two eigenvectors, and we have just shown that they must be linearly independent. So we have $\beta = \gamma = 0$*, which implies that* $\alpha = 0$ *as well. We can continue in this manner to show that any* k *eigenvectors with distinct eigenvalues are linearly independent.* ■

Corollary 160 *Hence,*

$$rank(V) \geq k = rank(v_1, ..., v_k)$$

where $V = [v_1 \ldots v_n]$ *is a* $n \times n$ *matrix of eigenvectors for all eigenvalues, and*

$$\dim\left(\sum_{j=1}^{k} \mathbb{M}_j\right) \geq k$$

where $\mathbb{M}_j$ *is the characteristic manifold corresponding to* λ_j*.*

Observe that

$$\sum_{j=1}^{k} \mathbb{M}_j \equiv \left\{ x \in \mathbb{R}^n \,|\, x = \textstyle\sum_{j=1}^{k} x_j, x_j \in \mathbb{M}_j \right\} = Span(V)$$

is simply the subspace of $\mathbb{R}^n$ that is spanned by all the eigenvectors of A. From this the geometric interpretation of the condition $\dim\left(\sum_{j=1}^{k} \mathbb{M}_j\right) \geq k$ should be clear: It simply means that, if A has k distinct eigenvalues, then its eigenvectors span at least (emphasis: at least) k of the n dimensions of $\mathbb{R}^n$.

Example 161 *Consider the matrix:*

$$A = \begin{bmatrix} 2 & -1 \\ -3 & 0 \end{bmatrix}$$

Its characteristic polynomial is

$$|A - \lambda I| = \lambda^2 - 2\lambda - 3 = (\lambda + 1)(\lambda - 3)$$

Hence, its eigenvalues are $\lambda_1 = -1$ and $\lambda_2 = 3$. To find the eigenvectors we have to solve $Av_1 = \lambda_1 v_1$ for v_1 and $Av_2 = \lambda_2 v_2$ for v_2. Let's find first an eigenvector $v_1 = \left[\begin{smallmatrix} v_{11} \\ v_{21} \end{smallmatrix}\right]$ corresponding to $\lambda_1 = -1$:

$$\begin{aligned} (A - \lambda_1 I)v_1 &= 0 \Leftrightarrow \\ \begin{bmatrix} 3 & -1 \\ -3 & 1 \end{bmatrix} \begin{bmatrix} v_{11} \\ v_{21} \end{bmatrix} &= \begin{pmatrix} 0 \\ 0 \end{pmatrix} \Leftrightarrow \\ 3v_{11} - v_{21} &= 0 \end{aligned}$$

Thus, any $v_1 = \left[\begin{smallmatrix} v_{11} \\ v_{21} \end{smallmatrix}\right]$ such that. $3v_{11} = v_{21} \neq 0$ is an eigenvector for $\lambda_1 = -1$, and vice versa, any eigenvector corresponding to λ_1 is of the form $v_1 = \left[\begin{smallmatrix} v_{11} \\ v_{21} \end{smallmatrix}\right]$ such that $3v_{11} = v_{21} \neq 0$. Now consider the eigenvector $v_2 = \left[\begin{smallmatrix} v_{12} \\ v_{22} \end{smallmatrix}\right]$ corresponding to $\lambda_2 = 3$:

$$\begin{aligned} (A - \lambda_2 I)v_2 &= 0 \Leftrightarrow \\ v_{12} + v_{22} &= 0 \end{aligned}$$

Hence, the eigenvectors corresponding to $\lambda_2 = 3$ are of the form $v_2 = \left[\begin{smallmatrix} v_{12} \\ v_{22} \end{smallmatrix}\right]$ such that. $v_{12} = -v_{22} \neq 0$. Finally, the manifolds spanned by these eigenvectors are:

$$\begin{aligned} \mathbb{M}_1 &= \left\{ x \in \mathbb{R}^2 \,|\, x = (\alpha, 3\alpha) \text{ for some } \alpha \in \mathbb{R} \right\} \\ \mathbb{M}_2 &= \left\{ x \in \mathbb{R}^2 \,|\, x = (\alpha, -\alpha) \text{ for some } \alpha \in \mathbb{R} \right\} \end{aligned}$$

As obvious, $\dim(\mathbb{M}_1) = \dim(\mathbb{M}_2) = 1$, *and further* $\mathbb{M}_1 \cap \mathbb{M}_2 = \{0\}$,$\mathbb{M}_1 + \mathbb{M}_1 = \mathbb{R}^2$, *implying that* $\{\mathbb{M}_1, \mathbb{M}_2\}$ *is a subspace partition for* $\mathbb{R}^2$.

For symmetric matrices we can say something stronger. Here we discuss the eigenvectors of distinct eigenvalues; a more general theorem follows below.

Theorem 162 *Suppose k ($k \leq n$) eigenvalues $\{\lambda_1, ..., \lambda_k\}$ of A are distinct with A symmetric, and take any corresponding eigenvectors $\{v_1, ..., v_k\}$, defined by $v_j \neq 0, Av_j = \lambda_j v_j$ for $j = 1, ..., k$. Then, $\{v_1, ..., v_k\}$ are orthogonal.*
Proof. *Suppose v is an eigenvector for λ and w is an eigenvector for μ. Then*

$$w'Av = \lambda w'v$$

and

$$w'Av = \mu w'v$$

■

(where the second equation is derived by taking the transpose of $Aw = \mu w$ and postmultiplying by v). Thus,

$$0 = (\lambda - \mu)w'v$$

implying that $w'v = 0$, or that w and v are orthogonal.

Remark: Finally, we notice that the computation of eigenvalues for diagonal or triangular matrices is trivial: These are given simply by the diagonal elements.

Lemma 163 *Let A be an $n \times n$ diagonal or triangular matrix with diagonal elements $\{a_{jj}\}$. Then its eigenvalues are $\lambda_j = a_{jj}$, all $j = 1, ..., n$.*

Exercise 164 *Here is how to work out the proof: Take an upper triangular matrix A, and form the matrix $C = \lambda I - A$. Notice that $C = \lambda I - A$ is upper triangular as well. Compute the determinant $\xi(\lambda) \equiv |\lambda I - A| = |C|$ inductively starting from the first row and going down: Show thereby that*

$$\begin{aligned}
|C| &= c_{11}|C_{11}| - c_{12}|C_{12}| + ... \pm c_{1n}|C_{1n}| = \\
&= (\lambda - a_{11})|C_{11}| + 0 = \\
&= ... = \\
&= (\lambda - a_{11})...(\lambda - a_{nn})
\end{aligned}$$

to conclude the proof.

Lemma 165 *If A is idempotent (defined by $AA = A$) then the eigenvalues of A are 0 or 1.*
Proof. *$Ax = \lambda x \Rightarrow Ax = AAx = \lambda Ax = \lambda^2 x$, so $\lambda^2 = \lambda$ which implies $\lambda = 0$ or $\lambda = 1$.* ■

6.2 Diagonalization and Canonical Form of a Matrix

Definition 166 *A matrix A is **diagonalizable** iff there exist an invertible matrix V such that $\Lambda \equiv V^{-1}AV$ is diagonal. And then Λ is the **canonical form** of A.*

It's a trivial exercise to show that:

$$\begin{aligned}
\Lambda &= V^{-1}AV \Rightarrow A = V\Lambda V^{-1} \Rightarrow \\
&\Rightarrow A - \lambda I = V\Lambda V^{-1} - \lambda V V^{-1} = V(\Lambda - \lambda I)V^{-1} \Rightarrow \\
&\Rightarrow |A - \lambda I| = |V|\,|\Lambda - \lambda I|\,|V^{-1}| = |\Lambda - \lambda I|
\end{aligned}$$

which means that:

Lemma 167 *If A is diagonalizable and Λ is its canonical, then A and Λ share the same characteristic polynomial and hence the same characteristic roots. And since Λ is diagonal, its eigenvalues are simply its diagonal elements. Thus, the canonical Λ of any matrix A, should it exists, is simply given by the eigenvalues $\{\lambda_j\}$ of matrix A.*

A natural question to make, Are all matrices diagonalizable? Unfortunately not all matrices are diagonalizable, but most[8] are.

Consider an $n \times n$ matrix A, and let $\{\lambda_1, ..., \lambda_n\}$ be its eigenvalues and $\{v_1, ..., v_n\}$ the corresponding eigenvectors. By definition then,

$$Av_j = \lambda_j v_j \quad \forall j = 1, ..., n \tag{9}$$

or in matrix form

$$[\,Av_1 \; Av_2 \; ... \; Av_n\,] = [\,\lambda_1 v_1 \; \lambda_2 v_2 \; ... \; \lambda_n v_n\,] \tag{10}$$

[Notice that both Av_j and $\lambda_j v_j$ are $n \times 1$ column vectors and hence the above matrices are $n \times n$.] Now form an $n \times n$ matrix V by stacking the eigenvectors v_j as columns,

$$V = [v_1 \, v_2 \, ... \, v_n] = \begin{bmatrix} v_{11} & v_{12} & ... & v_{1n} \\ v_{21} & v_{22} & ... & v_{2n} \\ ... & ... & & ... \\ v_{n1} & v_{n2} & ... & v_{nn} \end{bmatrix}$$

and a diagonal $n \times n$ matrix Λ with the eigenvalue λ_j as its (j, j) diagonal element,

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & ... & 0 \\ 0 & \lambda_2 & ... & 0 \\ ... & ... & & ... \\ 0 & 0 & ... & \lambda_n \end{bmatrix}$$

We can then rewrite (9) or (10) as follows:

$$AV = V\Lambda \tag{11}$$

[8] More precisely, given the set of all $n \times n$ matrices, the (sub)set of non-diagonalizable matrices is nonempty, but it is of measure zero. Indeed, if a matrix is nondiagonalizable, an arbitrarily small 'perturbartion' in its elements can make it diagonalizable, simply by making all its eigenvalues distinct.

So far we haven't assumed anything about the eigenvalues or the eigenvectors of A. It is always possible, given any A, to find a matrix V such that $AV = V\Lambda$. But the situation gets more interesting if V is invertible.

So, suppose now that A admits n linearly independent eigenvectors. This means that V is nonsingular and thus invertible. Thus, pre-multiplying both sides of (11) with V^{-1}, we get:

$$V^{-1}AV = \Lambda \tag{12}$$

or equivalently:

$$A = V\Lambda V^{-1} \tag{13}$$

The proposition works also in the other direction: Suppose that given a matrix A you can find an invertible matrix V and a diagonal matrix Λ such that (12) holds. Then any diagonal element of Λ is an eigenvalue for A, and the corresponding column of V is an eigenvector for that eigenvalue.

More precisely the following is true:

Theorem 168 *Let A be $n \times n$. If A has as many as n linearly independent eigenvectors, which means*

$$\begin{aligned} rank(V) &= \dim\left(\sum_{j=1}^{k} \mathbb{M}_j\right) = n \\ Span(V) &= \sum_{j=1}^{k} \mathbb{M}_j = \mathbb{R}^n \end{aligned}$$

then and only then A is diagonalizable.

Remark: The above result holds independently of whether the n eigenvalues are distinct or not.

We emphasize that for a matrix to be diagonalizable it is both necessary and sufficient that it admits n linearly independent eigenvectors. On the other hand, it is not necessary that it has n distinct eigenvalues. For instance, the $n \times n$ identity matrix I is trivially diagonalizable (for it is diagonal itself) but it has a single eigenvalue, the unit.

However, that A has n distinct eigenvalues is a sufficient condition for A to be diagonalizable:

Corollary 169 *Let A be $n \times n$. If A has n distinct eigenvalues, then it admits n linearly independent eigenvectors, and thus it is diagonalizable.*

Notice the geometric interpretation of $\sum_{j=1}^{k} \mathbb{M}_j = \mathbb{R}^n$: The eigenvectors $\{v_1, ..., v_k\}$ of A (for $v_j \in \mathbb{M}_j$) span the whole space $\mathbb{R}^n$, and $\{\mathbb{M}_1, ..., \mathbb{M}_k\}$ forms a partition of the whole $\mathbb{R}^n$.

Notice also that invertibility of a matrix A does not imply diagonalizability, as the following counterexample shows:

$A = \begin{bmatrix} 4 & 1 \\ -1 & 2 \end{bmatrix}$ is invertible. However, it has one eigenvalue (3) with multiplicity 2, and when we subtract 3 from each of its diagonal elements we find that

$$\begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = 0$$

is satisfied only when $v_1 = -v_2$. That is, all eigenvectors are multiples of $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$. There are not, then, 2 linearly independent eigenvectors for this matrix, and so this is an invertible matrix which is not diagonalizable.

But we can say something like the converse: if a matrix is diagonalizable, and if none of its eigenvalues are zero, then it is invertible. For in this case, we can invert Λ, and so we can write $A^{-1} = V^{-1}\Lambda^{-1}V$ (check that this is indeed the inverse of $A = V^{-1}\Lambda V$).

6.3 Symmetric Matrices

Observe also that the last theorem provides us with a sufficient and necessary condition, while the last lemma with only a sufficient condition. It is possible that a $n \times n$ matrix has n linearly independent vectors and is thus diagonalizable even when it does not have n distinct eigenvalues. To get a striking example of this, consider the $n \times n$ identity matrix I; notice that I has a single eigenvalue, the unit; yet, all $x \in \mathbb{R}^n$ $(x \neq 0)$ are eigenvectors, simply because $Ix = 1x$, and $\sum_j \mathbb{M}_j \equiv \{x \in \mathbb{R}^n | Ix = 1x\} = \mathbb{R}^n$.

More generally we have:

Lemma 170 *Let A be a **symmetric** $n \times n$ matrix. Then*
(i) All the eigenvalues and all the eigenvectors of A are real;
(ii) A admits n linearly independent eigenvectors and thus it is diagonalizable; and
*(iii) we can indeed find n eigenvectors v_j such that $v_j'v_j = 1$ and $v_i'v_j = 0$ $\forall i \neq j$. This means that $V = [v_j]$ is **orthonormal**, so that $V^{-1} = V'$ and A is diagonalizable as*

$$V'AV = \Lambda$$

That is, the so-constructed eigenvector matrix V is an orthonormal basis for $\mathbb{R}^n$ and the eigenspaces $\{\mathbb{M}_j | j = 1, ..., k\}$ form an orthogonal partition of $\mathbb{R}^n$.

Proof. *We will prove only part (iii), and thus (ii). Note that we have already shown these to be true for distinct eigenvalues, so our main concern here is repeated eigenvalues. Part (i) requires some complex analysis; if you are interested, look up Hermetian matrices in a linear algebra text (or google it) (question for future students - does that sound dated, or has the Google empire only expanded with time?); Hermetian matrices are the complex analog to symmetric matrices in the real field.*

(Step 1): Consider first a 2×2 symmetric matrix, and take as given that both its eigenvalues are real. Take any eigenvector v corresponding to the eigenvalue λ, so that $Av = \lambda v$. We may assume without loss of generality that $|v| = 1$

(this is merely a normalization). We claim that any w such that $w'v = 0$, and with $|w| = 1$ (again without loss of generality), is also an eigenvector. To see this, note that if A is symmetric, then $(Aw)'v = w'Av = w'\lambda v = \lambda w'v = 0$; that is, Aw is orthogonal to v. In $\mathbb{R}^2$, this means that Aw and w must point in the same direction, i.e. $Aw = \mu w$ for some μ. But this means that w is indeed an eigenvector, with eigenvalue μ.

(Step 2): We will now set up an argument which applies to any linear transformation mapping a two-dimensional vector space onto itself, $L : W \to W$, such that $(Lx)'y = x'(Ly)$. Note that a symmetric matrix satisfies this property.

Suppose we choose an arbitrary orthonormal basis $\{e_1, e_2\}$ for W. Now, since L maps W onto itself, we know that $Lx \in S(W)$ for any $x \in W$; in particular, it can be written as a linear combination of e_1 and e_2, since these are a basis. So write

$$Le_1 = b_{11}e_1 + b_{21}e_2, Le_2 = b_{12}e_1 + b_{22}e_2$$

where $b_{11}, b_{12}, b_{21}, b_{22}$ are scalars. We can collect this system as follows:

$$\begin{bmatrix} e_1 & e_2 \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} Le_1 & Le_2 \end{bmatrix}$$

(Check the dimensions on this equation to make sure you understand why everything is conformable - note that for now we are speaking of e_1 and e_2 as arbitrary elements of W, and L as a linear transformation, but everything still works if we make e_1 and e_2 vectors and L a matrix, as we will below). What we have done here is written a 2×2 matrix, $[b_{ij}] \equiv B$, to represent the linear transformation L. Now, because $\{e_1, e_2\}$ are orthonormal, and $(Lx)'y = x'(Ly)$ for $x, y \in W$, it is clear that $b_{21} = e_2'Le_1 = (Le_2)'e_1 = b_{12}$, so B is a symmetric 2×2 matrix.

(Step 3): Now we put the first two steps together. Consider a symmetric 3×3 matrix A with arbitrary (real) eigenvalues, possibly repeated. Suppose one eigenvalue/eigenvector pair satisfies $Av = \lambda v$, with $|v| = 1$. It is clear that v spans a one-dimensional subspace of $\mathbb{R}^3$; moreover, the nullspace of v spans a two-dimensional subspace $N(v) = \{x \in \mathbb{R}^3 : x'v = 0\}$. And with A symmetric, A maps $N(v)$ onto itself: $x \in N(v) \Rightarrow (Ax)'v = x'Av = x'\lambda v = \lambda x'v = 0 \Rightarrow Ax \in N(v)$. So A shares the qualities of the linear transformation L from Step 2. We will proceed as we did there, with our goal being to construct two more eigenvectors for A, both of which belong to $N(v)$ and which are also orthogonal to each other.

Choose an arbitrary orthonormal basis for $N(v)$ - this consists of two 3×1 vectors which we will again call $\{e_1, e_2\} \equiv E$. We can once again choose scalars b_{11}, b_{12}, b_{21}, and b_{22} such that

$$\begin{aligned} \begin{bmatrix} e_1 & e_2 \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} &= \begin{bmatrix} Ae_1 & Ae_2 \end{bmatrix} \\ \underset{(3\times 2)}{E}\underset{(2\times 2)}{B} &= \underset{(3\times 3)}{A}\underset{(3\times 2)}{E} \end{aligned}$$

In Step 1 we showed that any 2×2 symmetric real matrix admits a pair of orthonormal eigenvectors. So let $Bw_b = \mu w_b$ and $Bu_b = \nu u_b$, where w_b and u_b are orthonormal two-dimensional vectors and μ and ν are eigenvalues of B (not necessarily distinct). Note that $Ew_b = w$ and $Eu_b = u$ are 3×1 linear combinations of the columns of E, i.e. $w, u \in N(v)$. Then we have

$$\begin{aligned} Aw &= AEw_b = EBw_b = E\mu w_b = \mu E w_b = \mu w \\ Au &= AEu_b = EBu_b = E\nu u_b = \nu E u_b = \nu u \end{aligned}$$

So w and u are eigenvectors of A, with eigenvalues μ and ν, respectively. Moreover,

$$\begin{aligned} w'u &= (Ew_b)'Eu_b = w_b'E'Eu_b = w_b'I_2u_b = w_b'u_b = 0 \\ w'w &= (Ew_b)'Ew_b = w_b'E'Ew_b = w_b'I_2w_b = w_b'w_b = 1 \\ u'u &= (Eu_b)'Eu_b = u_b'E'Eu_b = u_b'I_2u_b = u_b'u_b = 1 \end{aligned}$$

so w and u are orthonormal. Since they are also orthonormal to the original eigenvector v, we have found an orthonomal set of eigenvectors for A; linear independence follows immediately.

The proof for general matrices A now follows by induction; we can continue following the above steps to show that any $n \times n$ symmetric matrix admits a set of n orthonormal eigenvectors. ■

So if A is symmetric, then, not only its eigenvectors span the whole space (that even when A is not itself a basis!), but we can also find a set of orthogonal eigenvectors, and thus the manifolds $\{\mathbb{M}_j\}$ are orthogonal to each other.

Exercise 171 *Consider the 2×2 identity matrix. What are its eigenvalues? Find a $V = [v_1 \ v_2]$ such that $V'V = I$ and $V^{-1}IV = I$. What are the corresponding $\{\mathbb{M}_1, \mathbb{M}_2\}$? Consider now*

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 3 \\ 0 & 3 & 1 \end{bmatrix}$$

Find an orthonormal V and a diagonal Λ such that $V'AV = \Lambda$.

6.4 Diagonalization as Changing Bases

Suppose we are given a linear transformation of the form

$$x \mapsto y = Ax$$

Now suppose that A is diagonalizable, and let V, Λ such that $V^{-1}AV = \Lambda$. Equivalently,

$$A = V\Lambda V^{-1}$$

implying

$$x \mapsto y = Ax = V\Lambda V^{-1}x$$

Now let us pre-multiply both sides of $y = V\Lambda V^{-1}x$ with V^{-1} to get

$$V^{-1}y = \Lambda V^{-1}x$$

Defining $\tilde{y} = (\tilde{y}_1, ..., \tilde{y}_n) = V^{-1}y$ and $\tilde{x} = (\tilde{x}_1, ..., \tilde{x}_n) = V^{-1}x$, we get

$$\tilde{y} = \Lambda\tilde{x}$$

or equivalently

$$\tilde{y}_i = \lambda_i \tilde{x}_i \quad \forall i = 1, ..., n$$

Since V^{-1} is invertible, the transformations

$$x \mapsto \tilde{x} = V^{-1}x \quad \text{and} \quad y \mapsto \tilde{y} = V^{-1}y$$

are one-to-one, and thus the transformation

$$x \mapsto y = Ax$$

is equivalent to the transformation

$$\tilde{x} \mapsto \tilde{y} = \Lambda\tilde{x}$$

or just to the system of simple transformations

$$\tilde{x}_i \mapsto \tilde{y}_i = \lambda_i \tilde{x}_i \quad \forall i = 1, ..., n$$

All that we did, in fact, is to change bases in $\mathbb{R}^n$. The initial basis was the customary one, the identity matrix I, and then x and y were the coordinates of two arbitrary vector elements with respect to the basis I. Now the new basis is V, and the new coordinates are $\tilde{x}$ and $\tilde{y}$. Indeed, notice that

$$x = V\tilde{x} \quad \text{and} \quad y = V\tilde{y}$$

which precisely means that $\tilde{x}$ and $\tilde{y}$ are the coordinates of x and y, respectively, with respect to the basis V. This should make clear the relevance of our discussion of how to change coordinates when we change bases. Recall that there we were considering two arbitrary bases E and F, and showed that the projection of E on F, or the matrix $P = F^{-1}E$, serves to transform the coordinates from E to F. In the present context, the initial basis is $E = I$ and the new basis is $F = V$, the matrix of n linearly independent eigenvectors, so that the change of coordinates is given by $x \mapsto Px$ for $P = F^{-1}E = V^{-1}$.

Finally, regarding the mapping $x \mapsto y = Ax$, observe the nice thing about the diagonalization of A and the corresponding change of bases: In the initial basis each value y_i depends on all $(x_1, ..., x_n)$, but in the new basis each $\tilde{y}_i$ depends only on the corresponding $\tilde{x}_i$ alone! The multidimensional transformation $x \mapsto y = Ax$ is equivalent to the single-dimensional transformations $\tilde{x}_i \mapsto \tilde{y}_i = \lambda_i \tilde{x}_i \ \ \forall i = 1, ..., n$, but the latter are much easier to work with.

The gains from this change of bases will indeed become clear when we apply the tools we developed here to solving dynamic systems, as well as in characterizing quadratic forms.