Introduction to Diagonalization

Suppose $T : \mathbb{R}^n \to \mathbb{R}^n$ is a linear transformation, and that $T(\mathbf{x}) = D\mathbf{x}$ where D is an $n \times n$ <u>diagonal</u> matrix. Such transformations are particularly easy to understand, and can be very useful in applications.

The effect of multiplication by D is just a rescaling in the directions of the basis coordinate vectors $e_1, e_2, ..., e_n$:

just rescales each basis vector: e_i maps to a multiple of itself, and the rescaling factor = the entry λ_i from the diagonal of D.

Since the \mathbf{e}_i 's are a basis, each vector $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = x_1 \mathbf{e}_1 + \dots + x_n \mathbf{e}_n$, so $T(\mathbf{x}) = D\mathbf{x} = D(x_1 \mathbf{e}_1 + \dots + x_n \mathbf{e}_n) = x_1 D(\mathbf{e}_1) + \dots + x_n D(\mathbf{e}_n)$ $= \lambda_1 x_1 \mathbf{e}_1 + \dots + \lambda_n x_n \mathbf{e}_n = \begin{bmatrix} \lambda_1 x_1 \\ \lambda_2 x_2 \\ \vdots \\ \lambda_n x_n \end{bmatrix}.$

In other words the effect of $T(\mathbf{x}) = D\mathbf{x}$ is just to rescale each coordinate x_i of \mathbf{x} by the corresponding factor λ_i .

<u>Example</u>: Suppose $T : \mathbb{R}^2 \to \mathbb{R}^2$ is given by $T(\boldsymbol{x}) = D\boldsymbol{x} = \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \boldsymbol{x}$.

Then $D\mathbf{e_1} = \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 3 \\ 0 \end{bmatrix} = 3\mathbf{e_1}, D\mathbf{e_2} = \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 6 \end{bmatrix} = 6\mathbf{e_2},$

and $D\boldsymbol{x} = \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \boldsymbol{x} = \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3x_1 \\ 6x_2 \end{bmatrix}.$

D simply rescales each coordinate of \boldsymbol{x} – by a factor of 3 in the direction of $\boldsymbol{e_1}$ (the horizontal direction) and by a factor of 6 in the direction of $\boldsymbol{e_2}$ (the vertical direction). In this particular example, the diagonal entries of *D* are both bigger than 1, so multiplying \boldsymbol{x} by *D* "stretches" the coordinates of \boldsymbol{x} in each of the coordinate directions, $\boldsymbol{e_1}$ and $\boldsymbol{e_2}$.

You can see this in the figure below where $\begin{bmatrix} 1\\1 \end{bmatrix} \mapsto D \begin{bmatrix} 1\\1 \end{bmatrix} = \begin{bmatrix} 3\\6 \end{bmatrix}$:



Example (the example is for \mathbb{R}^2 , but the same idea applies just as well for \mathbb{R}^n). Suppose

 $T(\boldsymbol{x}) = D\boldsymbol{x} = \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{2} \end{bmatrix} \boldsymbol{x}$. Consider the difference equation

$$\boldsymbol{x_{k+1}} = D\boldsymbol{x_k}$$

It's easy to see the long-term behavior of this system since multiplication by D just rescales coordinates: if we start with $\mathbf{x_0} = \begin{bmatrix} a \\ b \end{bmatrix}$, then $\mathbf{x_1} = D \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 1 \cdot a \\ \frac{1}{2} \cdot b \end{bmatrix}$, $\mathbf{x_2} = D\mathbf{x_1} = D \begin{bmatrix} 1 \cdot a \\ \frac{1}{2} \cdot b \end{bmatrix} = \begin{bmatrix} 1^2 \cdot a \\ (\frac{1}{2})^2 \cdot b \end{bmatrix}$

$$\boldsymbol{x_3} = D\boldsymbol{x_2} = D \begin{bmatrix} 1^2 \cdot a \\ (\frac{1}{2})^2 \cdot b \end{bmatrix} = \begin{bmatrix} 1^3 \cdot a \\ (\frac{1}{2})^3 \cdot b \end{bmatrix}$$
$$\boldsymbol{x_k} = \begin{bmatrix} 1^k \cdot a \\ (\frac{1}{2})^k \cdot b \end{bmatrix} = \begin{bmatrix} a \\ (\frac{1}{2})^k \cdot b \end{bmatrix}$$

and in general,

Since $(\frac{1}{2})^k \to 0$ as $k \to \infty$, we can immediately see the long term behavior of the system: $\boldsymbol{x}_k \to \begin{bmatrix} a \\ 0 \end{bmatrix}$ as $k \to \infty$.

In the same way, for any diagonal matrix $D = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$ and $\mathbf{x}_0 = \begin{bmatrix} a \\ b \end{bmatrix}$, we get $\mathbf{x}_k = \begin{bmatrix} \lambda_1^k a \\ \lambda_2^k b \end{bmatrix}$, so what happens to \mathbf{x}_k is determined by the size of

 λ_1, λ_2 : for example, if both $|\lambda_i| < 1$. then $\boldsymbol{x_k} \to \boldsymbol{0}$. The long term behavior of such a system can be very important in applications, and it depends on the diagonal entries of D.

Suppose $T : \mathbb{R}^n \to \mathbb{R}^n$, where $T(\boldsymbol{x}) = A\boldsymbol{x}$. If A is not diagonal, then <u>sometimes</u> it's possible to find a <u>new basis</u> \mathcal{B} so that a change to working in \mathcal{B} coordinates "diagonalizes" A. Roughly, this means that

(*) $\begin{cases} \boldsymbol{x} \mapsto A\boldsymbol{x} & \text{(in standard coordinates)} \quad \underline{\text{is the same as}} \\ [\boldsymbol{x}]_{\mathcal{B}} \mapsto D[\boldsymbol{x}]_{\mathcal{B}} & \text{(expressed instead in } \mathcal{B} \text{ coordinates)} \end{cases}$

To make this more precise:

For a basis \mathcal{B} , recall that $P_{\mathcal{B}}$ is the matrix whose columns are the new basis vectors in \mathcal{B} . Then standard coordinates and \mathcal{B} coordinates are related by the equations

 $P_{\mathcal{B}}[\boldsymbol{x}]_{\mathcal{B}} = \boldsymbol{x}$ and $[\boldsymbol{x}]_{\mathcal{B}} = P_{\mathcal{B}}^{-1}\boldsymbol{x}$

Sometimes we can find a basis \mathcal{B} and a diagonal matrix D so that if we:

1)	first convert \boldsymbol{x} to $\boldsymbol{\mathcal{B}}$ coordinates	$P_{\mathcal{B}}^{-1} oldsymbol{x}$
2)	then multiply by D (rescale \mathcal{B} coordinates)	$DP_{\mathcal{B}}^{-1}oldsymbol{x}$
3)	then convert back to standard coordinates	$P_{\mathcal{B}}DP_{\mathcal{B}}^{-1}\boldsymbol{x}$

we get the same result as $A\boldsymbol{x}$. In other words, $A\boldsymbol{x} = P_{\mathcal{B}}D P_{\mathcal{B}}^{-1}\boldsymbol{x}$

In that case, the effect of the mapping $\boldsymbol{x} \mapsto A\boldsymbol{x}$ is a <u>rescaling of the *B* coordinates</u> of \boldsymbol{x} with rescaling factors = the entries on the diagonal of *D*.

This leads to the official definition:

Suppose A is an $n \times n$ matrix. A is called <u>diagonalizable</u> if there is an invertible $n \times n$ matrix P and a diagonal matrix D such that

$$A = PDP^{-1}$$

(When this happens, the columns of P can be used as a new basis \mathcal{B} for \mathbb{R}^n and then $P = P_{\mathcal{B}}$ = the "change of coordinates matrix" so that things look just as they did in the preceding paragraphs.)

It is <u>not</u> always possible to diagonalize an $n \times n$ matrix A. We will see more about this in Chapter 5. But the examples below will lead us to one statement about when it is possible.

Example Suppose $T(\boldsymbol{x}) = A\boldsymbol{x}$, where $A = \begin{bmatrix} 4 & 2 \\ 1 & 5 \end{bmatrix}$. It's a fact (*just check by multiplying, if you like*) that

$$A = \begin{bmatrix} 4 & 2 \\ 1 & 5 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 3 & 0 \\ 0 & 6 \end{bmatrix} \begin{bmatrix} \frac{1}{3} & -\frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} \end{bmatrix}$$
$$= \begin{array}{c} \uparrow & \uparrow \\ P & D & P^{-1} \end{array}$$

(for now, don't worry about where the matrix P came from. Such questions come up in Chapter 5)

P is invertible, so its columns are linearly independent and span \mathbb{R}^2 and we can use the columns as a new basis for \mathbb{R}^2 : $\mathcal{B} = \{ \boldsymbol{b_1}, \boldsymbol{b_2} \} = \{ \begin{bmatrix} 2 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \}$. Then *P* is the matrix that the textbook calls $P_{\mathcal{B}} =$ "the change of coordinates matrix from \mathcal{B} coordinates to standard coordinates" :

$$P_{\mathcal{B}}[\boldsymbol{x}]_{\mathcal{B}} = \boldsymbol{x}$$

This new basis lets us see more clearly how A operates. Suppose $\boldsymbol{x} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$. Then

$$A\begin{bmatrix}2\\1\end{bmatrix} = \begin{bmatrix}4 & 2\\1 & 5\end{bmatrix}\begin{bmatrix}2\\1\end{bmatrix} = \begin{bmatrix}10\\7\end{bmatrix}$$

But we can also calculate Ax in a more roundabout way – but one which gives new insight:

$$A\begin{bmatrix}2\\1\end{bmatrix} = \begin{bmatrix}4&2\\1&5\end{bmatrix}\begin{bmatrix}2\\1\end{bmatrix} = \begin{bmatrix}2&1\\-1&1\end{bmatrix}\begin{bmatrix}3&0\\0&6\end{bmatrix}\begin{bmatrix}\frac{1}{3}&-\frac{1}{3}\\\frac{1}{3}&\frac{2}{3}\end{bmatrix}\begin{bmatrix}2\\1\end{bmatrix} = \begin{bmatrix}10\\7\end{bmatrix}$$
$$\stackrel{\uparrow}{A} \quad \stackrel{\uparrow}{\boldsymbol{x}} = P_{\mathcal{B}} \quad D \quad P_{\mathcal{B}}^{-1} \quad \stackrel{\uparrow}{\boldsymbol{x}}$$

In the second calculation, look at what happens to \boldsymbol{x} , step-by-step, as each matrix, in turn, does its work:

$$\boldsymbol{x} = \begin{bmatrix} 2\\1 \end{bmatrix} \mapsto \begin{bmatrix} \frac{1}{3} & -\frac{1}{3}\\ \frac{1}{3} & \frac{2}{3} \end{bmatrix} \begin{bmatrix} 2\\1 \end{bmatrix} = \begin{bmatrix} \frac{1}{3}\\ \frac{4}{3} \end{bmatrix}. \quad \underline{\text{Then}} \begin{bmatrix} 3 & 0\\0 & 6 \end{bmatrix} \begin{bmatrix} \frac{1}{3}\\ \frac{4}{3} \end{bmatrix} = \begin{bmatrix} 1\\8 \end{bmatrix}$$

multiplication by $\uparrow \qquad \uparrow \qquad \uparrow$
 P_B^{-1} converts $[\boldsymbol{x}]_{\mathcal{B}}$ the diagonal matrix stretched
standard coordinates
into \mathcal{B} -coordinates
 \mathcal{B} -coordinates by factors of 3 and 6
(in the coordinate directions
corresponding to \boldsymbol{b}_1 and \boldsymbol{b}_2)
 $\underline{\text{Finally}}, \qquad \begin{bmatrix} 2 & 1\\-1 & 1 \end{bmatrix} \begin{bmatrix} 1\\8 \end{bmatrix} = \begin{bmatrix} 10\\7 \end{bmatrix}$
 \uparrow
multiplication by P_B
converts the (stretched)
 \mathcal{B} -coordinates back into
standard coordinates

Use the figure on the following page to check each of the preceding steps graphically, as accurately as you can.

1)
$$\boldsymbol{x} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$
 has \mathcal{B} -coordinates $[\boldsymbol{x}]_{\mathcal{B}} = \begin{bmatrix} \frac{1}{3} \\ \frac{4}{3} \end{bmatrix}$

2) Stretching $[\boldsymbol{x}]_{\mathcal{B}}$ by factors of 3 and 6 in the $\boldsymbol{b_1}$ and $\boldsymbol{b_2}$ directions gives the point with "stretched" \mathcal{B} -coordinates $\begin{bmatrix} 1\\ 8 \end{bmatrix}$.

3) Converting these \mathcal{B} -coordinates $\begin{bmatrix} 1\\ 8 \end{bmatrix}$ back to standard coordinates gives $\begin{bmatrix} 10\\ 7 \end{bmatrix} = A \begin{bmatrix} 2\\ 1 \end{bmatrix}$.



So the effect of A on a point \boldsymbol{x} (in standard coordinates) is the same as the effect of the diagonal matrix D on the same point \boldsymbol{x} (when described in \mathcal{B} -coordinates) : a rescaling of the \mathcal{B} coordinates by a factor of 3 (in the direction of the "positive $\boldsymbol{b_1}$ -axis") and by a factor of 6 (in the direction of the "positive $\boldsymbol{b_2}$ -axis)."

Notice that
$$\boldsymbol{b_1} = \begin{bmatrix} 2\\ -1 \end{bmatrix}$$
 has \mathcal{B} coordinates $\begin{bmatrix} 1\\ 0 \end{bmatrix}$ and $\boldsymbol{b_2}$ has \mathcal{B} coordinates $\begin{bmatrix} 0\\ 1 \end{bmatrix}$

so (by the <u>underlined remark</u> in the preceding paragraph)

$$A\begin{bmatrix}2\\-1\end{bmatrix} = A\boldsymbol{b_1} = \text{the point with } \mathcal{B} \text{ coords.} \begin{bmatrix}3\\0\end{bmatrix} = 3\boldsymbol{b_1} = 3\begin{bmatrix}2\\-1\end{bmatrix} (\text{std coords.})$$

and

$$A\begin{bmatrix}1\\1\end{bmatrix} = A\mathbf{b_2} = \text{the point with } \mathcal{B} \text{ coords.} \begin{bmatrix}0\\6\end{bmatrix} = 6\mathbf{b_2} = 6\begin{bmatrix}1\\1\end{bmatrix} \text{ (in std coords)}$$

Again, it might be helpful to look at what's going on in more detail, using the same 3-step process. See the figure on the next page.

For example, for b_1 (follow each step using the preceding figure) :

Thus the matrix A maps each vector b_1 and b_2 to a scalar multiple of itself. Such (nonzero) vectors are called <u>eigenvectors</u> of A; the scalar 3 is called the <u>eigenvalue</u> associated with the eigenvector b_1 , and 6 is called the eigenvalue associated with the eigenvector b_2 .

The reason we were able to do such a nice analysis of how this matrix A works is that we were able to write down a new basis \mathcal{B} for \mathbb{R}^2 – <u>a basis whose members are eigenvectors</u> of A. (Remember, in the beginning, I simply <u>gave</u> you the matrix P whose columns $\mathbf{b_1 b_2}$ turned out to be eigenvectors; we haven't discussed yet how you might start, from "", to find such a)

Here is the general definition.

Definition A <u>nonzero</u> vector \boldsymbol{x} is called an <u>eigenvector</u> of the $n \times n$ matrix A if $A\boldsymbol{x} = \lambda \boldsymbol{x}$ for some scalar λ . The scalar λ is called an <u>eigenvalue</u> of A (associated with the eigenvector \boldsymbol{x}). Eigenvalues and eigenvectors of A are also called eigenvectors or eigenvalues of the transformation T if $T(\boldsymbol{x}) = A\boldsymbol{x}$,

It's traditional, in almost all books, to use the Greek letter λ ("lambda") to denote an eigenvalue.

<u>Eigenvalue</u> and <u>eigenvector</u> are words with German roots. Some books call "eigenvectors" and "eigenvalues" by the loose translations "characteristic vectors" and "characteristic values.")

In the preceding 2×2 example there is nothing special about the specific eigenvalues 3 and 6. They might be any other numbers λ_1 and λ_2 (where perhaps even $\lambda_1 = \lambda_2$). To recap, in general:

If A can be factored as $A = P\begin{bmatrix} \lambda_1 & 0\\ 0 & \lambda_2 \end{bmatrix} P^{-1}$, where P has columns $\boldsymbol{b_1}, \boldsymbol{b_2}$, then $\boldsymbol{b_1}$ and $\boldsymbol{b_2}$ are <u>eigenvectors</u> of A with <u>eigenvalues</u> λ_1 and λ_2 <u>because</u>:

We can choose $\mathcal{B} = \{b_1, b_2\}$ as a new <u>basis</u> (*why?*) for \mathbb{R}^2 , and then calculate

$$A\boldsymbol{b_1} = P_{\mathcal{B}} \begin{bmatrix} \lambda_1 & 0\\ 0 & \lambda_2 \end{bmatrix} P_{\mathcal{B}}^{-1} \boldsymbol{b_1} = P_{\mathcal{B}} \begin{bmatrix} \lambda_1 & 0\\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} 1\\ 0 \end{bmatrix}$$
$$|\underline{\qquad}|$$

$$\mathcal{B}\text{-coordinates of } \boldsymbol{b_1}$$

$$= P_{\mathcal{B}} \begin{bmatrix} \lambda_1 \\ 0 \end{bmatrix} = \begin{bmatrix} \mathbf{b_1} & \mathbf{b_2} \end{bmatrix} \begin{bmatrix} \lambda_1 \\ 0 \end{bmatrix} = \lambda_1 \mathbf{b_1} + 0 \mathbf{b_2} = \lambda_1 \mathbf{b_1}$$

$$|\underline{\qquad}|$$

$$\mathcal{B}\text{-coordinates}$$

Similarly, $A\boldsymbol{b_2} = \lambda_2 \boldsymbol{b_2}$.

In this case, the <u>geometric</u> interpretation of how A operates on a point \boldsymbol{x} is just as it was earlier: there is a rescaling of \mathcal{B} coordinates, with rescaling factors $= \lambda_1$ and λ_2 . Of course, if (say) $0 < \lambda_1 < 1$ then multiplication by D is a "contraction" of the first \mathcal{B} coordinate of \boldsymbol{x} (rather than a "stretch"); and if (say) $\lambda_2 < 0$, multiplication by D reverses the sign of the second \mathcal{B} coordinate of \boldsymbol{x} (as well as either stretching or contracting).

The preceding discussion proves the following theorem (where n = 2), and the discussion is exactly the same for $n \times n$ matrices (except that there are more matrix entries and vectors to write down, making the discuss <u>look</u> as if it were more complicated – it isn't !) So we have indicated the proof of the following theorem:

Theorem 1 Let A be an $n \times n$ matrix

Suppose A is diagonalizable, that is, suppose A can be factored as

	λ_1	0	0	0	0	
	0	λ_2	0	•••	0	
A = P	0		λ_3		0	P^{-1}
	0	÷	÷	÷	0	
	0	0	0	0	λ_n	

<u>Then</u> \mathbb{R}^n has a basis that consists of eigenvectors of A.

This basis is $\mathcal{B} = \{b_1, b_2, \dots, b_n\}$, where the b_i 's are the columns of P. These b_i 's are eigenvectors of A, and their eigenvalues are the scalars of the diagonal of D (in <u>same order</u>: λ_1 for b_1 , etc.).

It turns out that the <u>converse</u> of Theorem 1 is also true: if we start with a basis for \mathbb{R}^n that consists of eigenvectors of A, <u>then</u> A must be diagonalizable – that is, A can be factored as $A = PDP^{-1}$.

The next example illustrates why this works.

<u>Example</u> Suppose $A = \begin{bmatrix} 7 & 2 \\ -4 & 1 \end{bmatrix}$. Consider that basis $\mathcal{B} = \{ \begin{bmatrix} 1 \\ -1 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ -2 \end{bmatrix} \}$.

The vectors b_1 and b_2 are eigenvectors for A as we can easily check:

$$A\begin{bmatrix}1\\-1\end{bmatrix} = \begin{bmatrix}7 & 2\\-4 & 1\end{bmatrix}\begin{bmatrix}1\\-1\end{bmatrix} = \begin{bmatrix}5\\-5\end{bmatrix} = 5\begin{bmatrix}1\\-1\end{bmatrix}$$

so $\begin{bmatrix} 1\\ -1 \end{bmatrix}$ is an eigenvector with <u>eigenvalue</u> 5. Similarly, $A\begin{bmatrix} 1\\ -2 \end{bmatrix} = \begin{bmatrix} 7 & 2\\ -4 & 1 \end{bmatrix} \begin{bmatrix} 1\\ -2 \end{bmatrix} = \begin{bmatrix} 3\\ -6 \end{bmatrix} = 3 \begin{bmatrix} 1\\ -2 \end{bmatrix}$ so $\begin{bmatrix} 1\\ -2 \end{bmatrix}$ is an eigenvector with <u>eigenvalue</u> 3.

So $\mathcal{B} = \{\boldsymbol{b_1}, \boldsymbol{b_2}\}$ is a basis for \mathbb{R}^2 consisting of eigenvectors of A.

<u>Create</u> a matrix $P = [\boldsymbol{b_1} \ \boldsymbol{b_2}] = \begin{bmatrix} 1 & 1 \\ -1 & -2 \end{bmatrix}$. Because the columns of P are linearly independent, P is invertible.

<u>Create</u> $D = \begin{bmatrix} 5 & 0 \\ 0 & 3 \end{bmatrix}$, where the diagonal entries are the eigenvalues of b_1 , b_2 (written down in the same order: <u>not</u> $\begin{bmatrix} 3 & 0 \\ 0 & 5 \end{bmatrix}$!!)

Then $A = PDP^{-1}$. To see that this factorization works, the easiest thing to do (*rather than compute* P^{-1}) is to multiply both sides of the proposed equality on the right by P and check whether the equivalent equation AP = PD is true:

$AP = \left[\begin{array}{c} \\ -\end{array}\right]$	$\begin{bmatrix} 7 & 2 \\ 4 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -2 \end{bmatrix}$	$= \begin{bmatrix} 5\\-5 \end{bmatrix}$	$\begin{bmatrix} 3 \\ -6 \end{bmatrix}$ and
$PD = \begin{bmatrix} \\ - \end{bmatrix}$	$\begin{bmatrix} 1 & 1 \\ 1 & -2 \end{bmatrix} \begin{bmatrix} 5 \\ 0 \end{bmatrix}$	$\begin{bmatrix} 0\\3 \end{bmatrix}$	$= \begin{bmatrix} 5\\ -5 \end{bmatrix}$	$\begin{bmatrix} 3 \\ -6 \end{bmatrix}$

so the factorization is correct.

The general statement of the converse for Theorem 1 is

Theorem 2 Let A be an $n \times n$ matrix. <u>Suppose</u> \mathbb{R}^n <u>has a basis</u> { $\boldsymbol{b_1}, \boldsymbol{b_2}, ..., \boldsymbol{b_n}$ } <u>that consists of eigenvectors of A</u>. <u>Then A is diagonalizable</u>, that is $A = PDP^{-1}$ for some diagonal matrix D. The matrix P has the eigenvectors $\boldsymbol{b_1}, \boldsymbol{b_2}, ..., \boldsymbol{b_n}$ as its columns, and the diagonal matrix D has on its diagonal the corresponding eigenvalues $\lambda_1, \lambda_2, ..., \lambda_n$ (written in the same order: λ_1 for $\boldsymbol{b_1}$, etc.)

Why is Theorem 2 true? The argument is given here for the case n = 2, but it works in exactly the same way for \mathbb{R}^n . Not surprisingly, the argument in general looks almost exactly like the preceding example.

Suppose A is 2×2 and we know that $\mathcal{B} = \{\mathbf{b_1}, \mathbf{b_2}\}$ is a basis \mathbb{R}^2 consisting of eigenvectors of A. Let their corresponding eigenvalues be λ_1 and λ_2 .

Create matrix $P = \begin{bmatrix} \mathbf{b_1} & \mathbf{b_2} \end{bmatrix}$ and $D = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$. Since $\mathbf{b_1}$ and $\mathbf{b_2}$ are linearly independent, P is invertible. To complete the proof that A is diagonalizable, we that these matrices work: that $A = PDP^{-1}$.

As in the preceding example, we just need to verify that AP = PD, and, to do this, we simply need to remember the definition of matrix multiplication.

$$PD = P\begin{bmatrix} \lambda_1 & 0\\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} P\begin{bmatrix} \lambda_1\\ 0 \end{bmatrix} & P\begin{bmatrix} 0\\ \lambda_2 \end{bmatrix} \end{bmatrix}$$
$$= \begin{bmatrix} [\boldsymbol{b_1} \ \boldsymbol{b_2}] \cdot \begin{bmatrix} \lambda_1\\ 0 \end{bmatrix} & [\boldsymbol{b_1} \ \boldsymbol{b_2}] \cdot \begin{bmatrix} 0\\ \lambda_2 \end{bmatrix} = \begin{bmatrix} \lambda_1 \boldsymbol{b_1} \ \lambda_2 \boldsymbol{b_2} \end{bmatrix}, \text{ and}$$

$$AP = A[\mathbf{b_1} \ \mathbf{b_2}] = [A\mathbf{b_1} \ A\mathbf{b_2}] = [\lambda_1\mathbf{b_1} \ \lambda_2\mathbf{b_2}]$$

$$\uparrow$$

<u>because</u> b_1 is an eigenvector of A with eigenvalue λ_1 , and b_2 is an eigenvector of A with eigenvalue λ_2 .

Therefore AP and PD have the same columns, so AP = PD

<u>Question</u> Based on this discussion, can you give an example of a transformation $T : \mathbb{R}^2 \to \mathbb{R}^2$ where $T(\mathbf{x}) = A\mathbf{x}$ and A is <u>not</u> diagonalizable?

The Big Picture, so far

In Chapter 4, we have been discussing vector spaces V (where V might not be \mathbb{R}^n). After discussing linear independence and spanning in this more general setting, we were led to the idea of a basis $\mathcal{B} = \{\mathbf{b_1}, ..., \mathbf{b_n}\}$ for V. From that, the Unique Representation Theorem (p. 246) led to the idea of using \mathcal{B} to create <u>coordinates</u> for each \mathbf{x} in V. The <u>coordinate vector</u> $[\mathbf{x}]_{\mathcal{B}}$ is a vector in \mathbb{R}^n .

We saw that the mapping $\boldsymbol{x} \mapsto [\boldsymbol{x}]_{\mathcal{B}}$ is an <u>isomorphism</u> (a one-to-one, onto linear mapping) from the vector space V to \mathbb{R}^n . This association preserves all vector space operations and all linear dependency relations. For example, a linearly independent set of <u>vectors</u> in V has a linearly independent set of <u>coordinate vectors</u> in \mathbb{R}^n , and vice-versa.

In the special case when $V = \mathbb{R}^n$, we can might choose a basis $\mathcal{B} = \{\boldsymbol{b_1}, ..., \boldsymbol{b_n}\}$ different from the standard basis. Then a vector \boldsymbol{x} in \mathbb{R}^n gets new "nonstandard" coordinates $[\boldsymbol{x}]_{\mathcal{B}}$ relative to the basis \mathcal{B} . The matrix $P_{\mathcal{B}} = [\boldsymbol{b_1} ... \boldsymbol{b_n}]$ is the "operator" that changes \mathcal{B} coordinates into standard coordinates according to the formula $P_{\mathcal{B}}[\boldsymbol{x}]_{\mathcal{B}} = \boldsymbol{x}$.

Since the columns of $P_{\mathcal{B}}$ are linearly independent, the change of coordinates matrix $P_{\mathcal{B}}$ is always invertible and $P_{\mathcal{B}}^{-1}$ is the "operator" that converts standard coordinates into \mathcal{B} -coordinates: $[\boldsymbol{x}]_{\mathcal{B}} = P_{\mathcal{B}}^{-1}\boldsymbol{x}$.

If a 2×2 matrix A can be factored into the form PDP^{-1} , where D is a diagonal matrix, then A is called <u>diagonalizable</u> – because in that case A "acts like a diagonal matrix" when computations are done relative to the basis B that consists of the columns. This led into the idea of eigenvectors and eigenvalues of the matrix A.

We proved that a 2×2 matrix A is diagonalizable <u>if and only if</u> \mathbb{R}^2 has a basis consisting of eigenvectors of A – and indicated that a completely similar proof works for an $n \times n$ matrix A operating on \mathbb{R}^n .

The conceptual idea of diagonalization and its relation to a basis of eigenvectors is nicely motivated geometrically and not very hard. You may have noticed, however, that in the preceding examples:

a factorization of a given matrix A into PDP^{-1} was given, or

the eigenvalues and eigenvectors of A were <u>given</u> so that P and D could be created.

But if you are simply given an $n \times n$ matrix A, then trying to <u>find</u> its eigenvectors and eigenvalues (*if it has any at all*!), and determining whether \mathbb{R}^n does or does not have a <u>basis</u> <u>consisting of eigenvectors</u> of A are harder questions. Be aware of those questions, but try to keep worries about them suppressed until we get into Chapter 5. For now just focus on the concept of diagonalization, what it means, and how diagonalization is connected to eigenvectors and eigenvalues.