Matrix Methods for Computational Modeling and Data Analytics

Virginia Tech · Spring 2019

Mark Embree

embree@vt.edu

version of 11 February 2019

Chapter 5
Orthogonality

WE NEED ONE MORE TOOL from basic linear algebra to tackle the applications that lie ahead. In the last chapter we saw that basis vectors provide the most economical way to describe a subspace, but not all bases are created equal. Clearly the following three pairs of vectors all form bases for $\mathbb{R}^2$:

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \begin{bmatrix} 0 \\ 1 \end{bmatrix};$$

$$\begin{bmatrix} \sqrt{2}/2 \\ \sqrt{2}/2 \end{bmatrix}, \qquad \begin{bmatrix} \sqrt{2}/2 \\ -\sqrt{2}/2 \end{bmatrix};$$

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \begin{bmatrix} 1, \\ .001 \end{bmatrix}.$$

For example, the vector $\mathbf{x} = [1, 1]^T$ can be written as a linear combination of all three sets of vectors. In the first two cases,

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} = 1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \sqrt{2} \begin{bmatrix} \sqrt{2}/2 \\ \sqrt{2}/2 \end{bmatrix} + 0 \begin{bmatrix} \sqrt{2}/2 \\ -\sqrt{2}/2 \end{bmatrix},$$

the coefficients multiplying against the basis vectors are no bigger than $\|\mathbf{x}\| = \sqrt{2}$. In the third case,

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} = -999 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 1000 \begin{bmatrix} 1 \\ .001 \end{bmatrix},$$

the coefficients -999 and 1000 are *much bigger* than $\|\mathbf{x}\|$. Small changes in $\mathbf{x}$ require significant swings in these coefficients. For example, if we change the second entry of $\mathbf{x}$ from 1 to 0.9,

$$\begin{bmatrix} 1 \\ .9 \end{bmatrix} = -899 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 900 \begin{bmatrix} 1 \\ .001 \end{bmatrix},$$

the coefficients change by 100, which is 1000 times the change in $\mathbf{x}$!

The first two bases above are special, for the norm of each vector is one, and the vectors are orthogonal. The vectors in the third basis form a small angle. Indeed, all these bases are linearly independent, but, extrapolating from *Animal Farm*, some bases are more linearly independent than others.

This chapter is devoted to the art of orthogonalizing a set of linearly independent vectors. The main tool of this craft is the orthogonal projector.

5.1 Projectors

Among the most fundamental operations in linear algebra is the *projection* of a vector into a particular subspace. Such actions are encoded in a special class of matrices.

Definition 18 *A matrix* $\mathbf{P} \in \mathbb{R}^{n\times n}$ *is a* projector *provided* $\mathbf{P}^2 = \mathbf{P}$. *If additionally* $\mathbf{P} = \mathbf{P}^T$, *then* $\mathbf{P}$ *is called an* orthogonal projector.

The simplest projectors send all vectors into a one-dimensional subspace. Suppose that $\mathbf{w}^T\mathbf{v} = 1$. Then

$$\mathbf{P} := \mathbf{v}\mathbf{w}^T$$

is a projector, since $\mathbf{P}^2 = (\mathbf{v}\mathbf{w}^T)(\mathbf{v}\mathbf{w}^T) = \mathbf{v}(\mathbf{w}^T\mathbf{v})\mathbf{w}^T = \mathbf{v}\mathbf{w}^T = \mathbf{P}$. We say that $\mathbf{P}$ projects *onto* $\operatorname{span}\{\mathbf{v}\}$, since

$$\mathcal{R}(\mathbf{P}) = \operatorname{span}\{\mathbf{v}\},$$

and *along* $\operatorname{span}\{\mathbf{w}\}^\perp$, since

$$\begin{aligned}\mathcal{N}(\mathbf{P}) &= \operatorname{span}\{\mathbf{w}\}^\perp \\ &= \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^T\mathbf{w} = 0\}.\end{aligned}$$

Example 1 *Consider the vectors*

$$\mathbf{v} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \mathbf{w} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Confirm that $\mathbf{w}^T\mathbf{v} = 1$, *so*

$$\mathbf{P} = \mathbf{v}\mathbf{w}^T = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$$

is a projector. It sends the vector

$$\mathbf{x} = \begin{bmatrix} 2 \\ 1 \end{bmatrix},$$

to

$$\mathbf{Px} = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 3 \\ 0 \end{bmatrix},$$

which is contained in $\mathcal{R}(\mathbf{P}) = \mathrm{span}(\mathbf{v})$. *The remainder,*

$$\mathbf{x} - \mathbf{Px} = \begin{bmatrix} 2 \\ 1 \end{bmatrix} - \begin{bmatrix} 3 \\ 0 \end{bmatrix} = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

is indeed orthogonal to $\mathrm{span}(\mathbf{w})$.

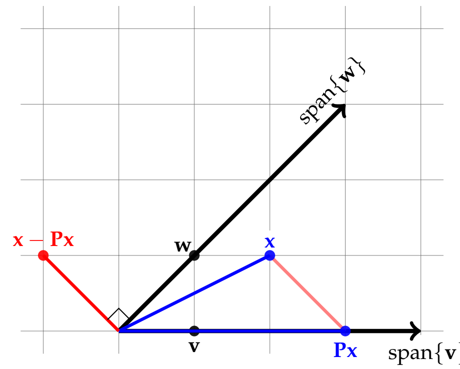

The projector $\mathbf{P}$ sends the vector $\mathbf{x}$ to $\mathbf{Px} \in \mathcal{R}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\}$, while the remainder $\mathbf{x} - \mathbf{Px}$ is in $\mathcal{N}(\mathbf{P}) = \mathrm{span}\{\mathbf{w}\}^{\perp}$, i.e., $\mathbf{x} - \mathbf{Px}$ is orthogonal to $\mathbf{w}$.

The remainder $\mathbf{x} - \mathbf{Px}$ plays a key role in our development. Notice that $\mathbf{P}$ annihilates this vector,

$$\mathbf{P}(\mathbf{x} - \mathbf{Px}) = \mathbf{Px} - \mathbf{P}^2\mathbf{x} = \mathbf{Px} - \mathbf{Px} = \mathbf{0},$$

which means that

$$\mathbf{x} - \mathbf{Px} \in \mathcal{N}(\mathbf{P}).$$

For the projector in the last example, $\mathbf{P} \neq \mathbf{P}^T$, so we call $\mathbf{P}$ an *oblique* projector. For such projectors, the remainder $\mathbf{x} - \mathbf{Px}$ generally forms an oblique angle with $\mathcal{R}(\mathbf{P})$. In this course we shall most often work with *orthogonal* projectors, for which $\mathbf{P} = \mathbf{P}^T$. Of course this implies $\mathcal{N}(\mathbf{P}) = \mathcal{N}(\mathbf{P}^T)$, and so, the Fundamental Theorem of Linear Algebra's $\mathcal{N}(\mathbf{P}^T) \perp \mathcal{R}(\mathbf{P})$ becomes

oblique = not 90°

$$\mathcal{N}(\mathbf{P}) \perp \mathcal{R}(\mathbf{P}).$$

Since $\mathbf{x} - \mathbf{Px} \in \mathcal{N}(\mathbf{P})$, the remainder is always orthogonal to $\mathcal{R}(\mathbf{P})$. This explains why we call such a $\mathbf{P}$ an *orthogonal* projector.

To build an orthogonal projector onto a one-dimensional subspace, we need only a single vector $\mathbf{v}$. Set

$$\mathbf{P} = \mathbf{v}\mathbf{v}^T,$$

and, to give $\mathbf{P}^2 = \mathbf{P}$, require that $\mathbf{v}^T\mathbf{v} = \|\mathbf{v}\|^2 = 1$, i.e., $\mathbf{v}$ must be a unit vector. This $\mathbf{P}$ projects *onto*

$$\mathcal{R}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\},$$

and *along* $\mathrm{span}\{\mathbf{v}\}$, since

$$\mathcal{N}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\}^{\perp}.$$

The notation $\mathrm{span}\{\mathbf{v}\}^{\perp}$ denotes the set of all vectors orthogonal to $\mathrm{span}\{\mathbf{v}\}$.

We illustrate these ideas in the following example.

Example 2 *Use the unit vector* $\mathbf{v}$ *from the last example to build*

$$\mathbf{P} = \mathbf{v}\mathbf{v}^T = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

which projects vectors onto the "x axis". For the same $\mathbf{x}$ *used before, we have*

$$\mathbf{Px} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \end{bmatrix}.$$

Now observe the orthogonality of the remainder

$$\mathbf{x} - \mathbf{Px} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

with the projection subspace $\mathcal{R}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\}$.

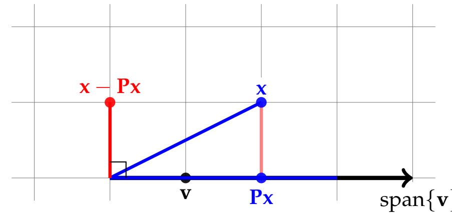

The orthogonal projector $\mathbf{P}$ sends the vector $\mathbf{x}$ to $\mathbf{Px} \in \mathcal{R}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\}$, while the remainder $\mathbf{x} - \mathbf{Px}$ is contained in $\mathcal{N}(\mathbf{P}) = \mathrm{span}\{\mathbf{v}\}^{\perp}$, i.e., $\mathbf{x} - \mathbf{Px}$ is orthogonal to $\mathbf{v}$ and $\mathbf{Px}$.

Did you notice that, in contrast to the first example, the projection $\mathbf{Px}$ has smaller norm than $\mathbf{x}$, $\|\mathbf{Px}\| \le \|\mathbf{x}\|$? Such reduction is a general trait of orthogonal projectors.

Proposition 2 *If* $\mathbf{P}$ *is an orthogonal projector, then* $\|\mathbf{Px}\| \le \|\mathbf{x}\|$ *for all* $\mathbf{x} \in \mathbb{R}^n$.

Proof. Since $\mathbf{P}$ is an orthogonal projector, $\mathbf{P}^T = \mathbf{P}$ and $\mathbf{P}^2 = \mathbf{P}$. Hence $\mathbf{Px}$ is orthogonal to $\mathbf{x} - \mathbf{Px}$:

$$\begin{aligned}(\mathbf{Px})^T(\mathbf{x} - \mathbf{Px}) &= \mathbf{x}^T\mathbf{P}^T\mathbf{x} - \mathbf{x}^T\mathbf{P}^T\mathbf{Px} \\ &= \mathbf{x}^T\mathbf{Px} - \mathbf{x}^T\mathbf{P}^2\mathbf{x} = \mathbf{x}^T\mathbf{Px} - \mathbf{xPx} = 0.\end{aligned}$$

Now apply the Pythagorean Theorem to the orthogonal pieces $\mathbf{Px} \in \mathcal{R}(\mathbf{P})$ and $\mathbf{x} - \mathbf{Px} \in \mathcal{N}(\mathbf{P})$,

$$\|\mathbf{x}\|^2 = \|\mathbf{Px} + (\mathbf{x} - \mathbf{Px})\|^2 = \|\mathbf{Px}\|^2 + \|\mathbf{x} - \mathbf{Px}\|^2,$$

and conclude the result by noting that $\|\mathbf{x} - \mathbf{Px}\| \ge 0$. ■

STUDENT EXPERIMENTS

5.18. To get an appreciation for projectors, in MATLAB construct the orthogonal projector $\mathbf{P}$ onto

$$\mathrm{span}\left\{\begin{bmatrix} 2 \\ 1 \end{bmatrix}\right\}.$$

Now construct 500 random vectors $\mathbf{x}$ = `randn(2,1)`. Produce a figure showing all 500 of these vectors in the plane (plotted as blue dots), then superimpose on the plot the values of $\mathbf{Px}$ (plotted as red dots). Use `axis equals` to scale the axes in the same manner.

Repeat the experiment, but now in three dimensions, with $\mathbf{P}$ the orthogonal projector onto

$$\mathrm{span}\left\{\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}\right\}.$$

Use MATLAB's `plot3` command plot $\mathbf{x}$ and $\mathbf{Px}$ in three dimensional space.

5.2 Orthogonalization

Now we arrive at the main point of this chapter: the transformation of basis into an orthonormal basis for the same set.

Definition 19 *A basis* $\{\mathbf{q}_1, \ldots, \mathbf{q}_n\}$ *is* orthonormal *provided*

- $\mathbf{q}_j^T \mathbf{q}_k = 0$ *for all* $j \neq k$ *(the vectors are orthogonal);*
- $\|\mathbf{q}_j\| = 1$ *for* $j = 1, \ldots, n$ *(the vectors are normalized).*

We can summarize orthonormality quite neatly if we arrange the basis in the matrix

$$\mathbf{Q} = [\, \mathbf{q}_1 \quad \mathbf{q}_2 \quad \cdots \quad \mathbf{q}_n \,].$$

Then $\mathbf{Q}^T\mathbf{Q}$ contains all possible inner products between $\mathbf{q}_j$ and $\mathbf{q}_k$:

$$\mathbf{Q}^T\mathbf{Q} = \begin{bmatrix} \mathbf{q}_1^T\mathbf{q}_1 & \mathbf{q}_1^T\mathbf{q}_2 & \cdots & \mathbf{q}_1^T\mathbf{q}_n \\ \mathbf{q}_2^T\mathbf{q}_1 & \mathbf{q}_2^T\mathbf{q}_2 & \cdots & \mathbf{q}_2^T\mathbf{q}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{q}_n^T\mathbf{q}_1 & \mathbf{q}_n^T\mathbf{q}_2 & \cdots & \mathbf{q}_n^T\mathbf{q}_n \end{bmatrix} = \mathbf{I},$$

since $\mathbf{q}_j^T\mathbf{q}_j = \|\mathbf{q}_j\|^2 = 1$.

Definition 20 *A matrix* $\mathbf{Q} \in \mathbb{R}^{n\times n}$ *is* unitary *if* $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$.
If $\mathbf{Q} \in \mathbb{R}^{m\times n}$ *with* $m > n$ *and* $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$, *then* $\mathbf{Q}$ *is* subunitary.

STUDENT EXPERIMENTS

5.19. Suppose $\mathbf{Q} \in \mathbb{R}^{m\times n}$ with $m < n$. Is it possible that $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$? Explain.

5.20. Suppose that $\mathbf{Q} \in \mathbb{R}^{m\times n}$ is subunitary. Show that $\mathbf{Q}\mathbf{Q}^T$ is a projector. What is $\mathcal{R}(\mathbf{Q})$ in terms of $\mathbf{q}_1, \ldots, \mathbf{q}_n$?

We are now prepared to orthogonalize a set of vectors. Suppose we have a basis $\mathbf{a}_1, \ldots, \mathbf{a}_n$ for a subspace $\mathcal{S} \subset \mathbb{C}^m$. We seek an orthonormal basis $\mathbf{q}_1, \ldots, \mathbf{q}_n$ for the same subspace, which we shall build one vector at a time. The goal is to first construct unit vector $\mathbf{q}_1$ so that

$$\text{span}\{\mathbf{q}_1\} = \text{span}\{\mathbf{a}_1\}$$

and then to build unit vector $\mathbf{q}_2$ orthogonal to $\mathbf{q}_1$ so that

$$\text{span}\{\mathbf{q}_1, \mathbf{q}_2\} = \text{span}\{\mathbf{a}_1, \mathbf{a}_2\}$$

and then unit vector $\mathbf{q}_3$ orthogonal to $\mathbf{q}_1$ and $\mathbf{q}_2$ so that

$$\text{span}\{\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3\} = \text{span}\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$$

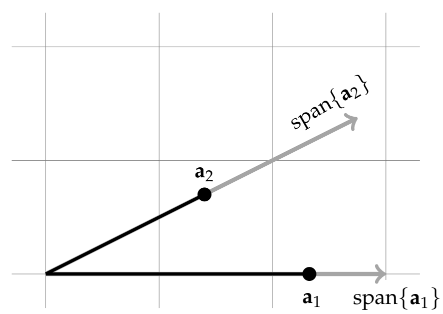

The setting: basis vectors $\mathbf{a}_1$ and $\mathbf{a}_2$ that are neither unit length, nor orthogonal.

and so on, one $\mathbf{q}_j$ vector at a time, until we have the whole new basis:

$$\operatorname{span}\{\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n\} = \operatorname{span}\{\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n\} = \mathcal{S}.$$

The first of these steps is easy, for we get $\mathbf{q}_1$ by normalizing $\mathbf{a}_1$:

$$\mathbf{q}_1 = \frac{1}{\|\mathbf{a}_1\|}\mathbf{a}_1.$$

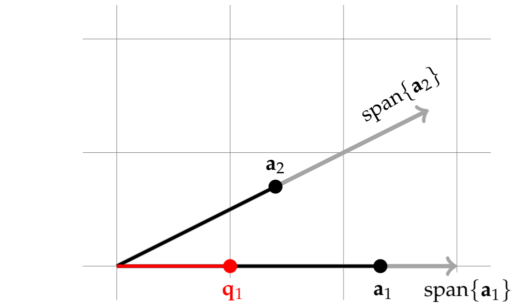

First normalize $\mathbf{a}_1$ to obtain $\mathbf{q}_1$ with $\operatorname{span}\{\mathbf{q}_1\} = \operatorname{span}\{\mathbf{a}_1\}$.

The next step is decisive, and here we use insight gained from our quick study of orthogonal projectors. Figuratively speaking, we must remove the part of $\mathbf{a}_2$ in the direction $\mathbf{q}_1$, leaving the portion of $\mathbf{a}_2$ orthogonal to $\mathbf{q}_1$. Since $\mathbf{q}_1$ is a unit vector,

$$\mathbf{P}_1 := \mathbf{q}_1\mathbf{q}_1^T$$

is an orthogonal projector onto $\operatorname{span}\{\mathbf{q}_1\} = \operatorname{span}\{\mathbf{a}_1\}$, so

$$\mathbf{P}_1\mathbf{a}_2 = \mathbf{q}_1\mathbf{q}_1^T\mathbf{a}_2$$

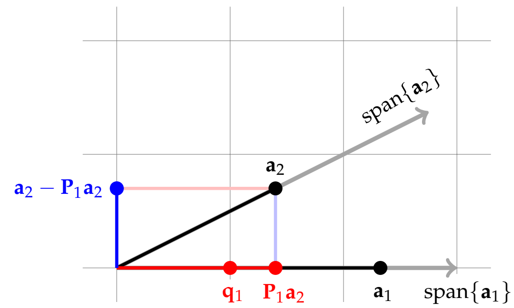

Compute $\widehat{\mathbf{q}_2} = \mathbf{a}_2 - \mathbf{P}_1\mathbf{a}_2$, the portion of $\mathbf{a}_2$ that is orthogonal to $\mathbf{q}_1$.

is the part of $\mathbf{a}_2$ in the direction $\mathbf{q}_1$. Remove that from $\mathbf{a}_2$ to get

$$\widehat{\mathbf{q}}_2 := \mathbf{a}_2 - \mathbf{P}_1\mathbf{a}_2.$$

And since

$$\mathbf{P}_1\widehat{\mathbf{q}}_2 = \mathbf{P}_1(\mathbf{a}_2 - \mathbf{P}_1\mathbf{a}_2) = \mathbf{0},$$

spot that $\widehat{\mathbf{q}}_2 \in \mathcal{N}(\mathbf{P}_1)$, which is orthogonal to $\mathcal{R}(\mathbf{P}_1) = \operatorname{span}\{\mathbf{q}_1\}$, so $\widehat{\mathbf{q}}_2 \perp \mathbf{q}_1$. If you prefer a less high-falutin' explanation, just compute

$$\mathbf{q}_1^T\widehat{\mathbf{q}}_2 = \mathbf{q}_1^T\mathbf{a}_2 - \mathbf{q}_1^T\mathbf{q}_1\mathbf{q}_1^T\mathbf{a}_2 = \mathbf{q}_1^T\mathbf{a}_2 - \mathbf{q}_1^T\mathbf{a}_2 = 0.$$

So we have constructed a vector $\widehat{\mathbf{q}}_2$ orthogonal to $\mathbf{q}_1$ such that

$$\operatorname{span}\{\mathbf{q}_1, \widehat{\mathbf{q}}_2\} = \operatorname{span}\{\mathbf{a}_1, \mathbf{a}_2\}.$$

Since we want not just an orthogonal basis, but an *orthonormal* basis, we adjust $\widehat{\mathbf{q}}_2$ by scaling it to become the unit vector

$$\mathbf{q}_2 := \frac{1}{\|\widehat{\mathbf{q}}_2\|}\widehat{\mathbf{q}}_2.$$

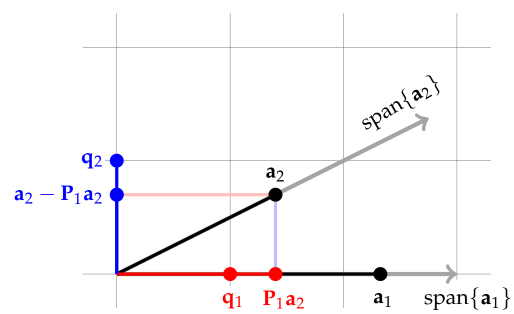

Normalize $\mathbf{q}_2 = \widehat{\mathbf{q}_2}/\|\widehat{\mathbf{q}_2}\|$ to get a unit vector in the $\widehat{\mathbf{q}_2}$ direction.

The orthogonalization process for subsequent vectors follows the same template. To construct $\mathbf{q}_3$, we first remove from $\mathbf{a}_3$ its components in the $\mathbf{q}_1$ and $\mathbf{q}_2$ directions:

$$\widehat{\mathbf{q}}_3 := \mathbf{a}_3 - \mathbf{P}_1\mathbf{a}_3 - \mathbf{P}_2\mathbf{a}_3,$$

where $\mathbf{P}_2 = \mathbf{q}_2\mathbf{q}_2^T$, where for $j = 1, 2$,

$$\begin{aligned}\mathbf{q}_j^T\widehat{\mathbf{q}}_3 &= \mathbf{q}_j^T\mathbf{a}_3 - \mathbf{q}_j^T\mathbf{q}_1\mathbf{q}_1^T\mathbf{a}_3 - \mathbf{q}_j^T\mathbf{q}_2\mathbf{q}_2^T\mathbf{a}_3 \\ &= \mathbf{q}_j^T\mathbf{a}_3 - \mathbf{q}_j^T\mathbf{a}_3 = 0,\end{aligned}$$

using the orthonormality of $\mathbf{q}_1$ and $\mathbf{q}_2$. Normalize $\widehat{\mathbf{q}}_3$ to get

$$\mathbf{q}_3 := \frac{1}{\|\widehat{\mathbf{q}}_3\|}\widehat{\mathbf{q}}_3.$$

Now the general pattern should be evident. For future vectors, construct

$$\widehat{\mathbf{q}}_{k+1} = \mathbf{a}_{k+1} - \sum_{j=1}^{k} \mathbf{P}_j \mathbf{a}_{k+1}, \qquad \text{with } \mathbf{P}_j := \mathbf{q}_j \mathbf{q}_j^T,$$

then normalize

$$\mathbf{q}_{k+1} = \frac{1}{\|\widehat{\mathbf{q}}_{k+1}\|}\widehat{\mathbf{q}}_{k+1},$$

giving

$$\mathbf{q}_{k+1} \perp \operatorname{span}\{\mathbf{q}_1, \ldots, \mathbf{q}_k\}$$

which extends the orthonormal basis in one more direction:

$$\operatorname{span}\{\mathbf{q}_1, \ldots, \mathbf{q}_k, \mathbf{q}_{k+1}\} = \operatorname{span}\{\mathbf{a}_1, \ldots, \mathbf{a}_k, \mathbf{a}_{k+1}\}.$$

This algorithm is known as the *Gram–Schmidt process*.

STUDENT EXPERIMENTS

5.21. Under what circumstances will this procedure *break down*? That is, will you ever divide by zero when trying to construct $\mathbf{q}_{k+1} = \widehat{\mathbf{q}}_{k+1} / \|\widehat{\mathbf{q}}_{k+1}\|$?

5.22. Show that $\mathbf{\Pi}_k = \mathbf{P}_1 + \cdots + \mathbf{P}_k$ is an orthogonal projector. With this notation, explain how we can compactly summarize each Gram–Schmidt step as

$$\mathbf{q}_{k+1} = \frac{(\mathbf{I} - \mathbf{\Pi}_k)\mathbf{a}_{k+1}}{\|(\mathbf{I} - \mathbf{\Pi}_k)\mathbf{a}_{k+1}\|}.$$

5.3 Gram–Schmidt is QR factorization

The Gram–Schmidt process is a classical way to orthogonalize a basis, and you can execute the process by hand when the vectors are short (m is small) and there are few of them (n is small). As with Gaussian elimination, we want to automate the process for larger problems – and just as Gaussian elimination gave us the $\mathbf{A} = \mathbf{LU}$ factorization of a square matrix, so Gram–Schmidt will lead to a factorization of a general rectangular matrix.

Let us work through the arithmetic behind the orthogonalization of three linearly independent vectors $\mathbf{a}_1$, $\mathbf{a}_2$, and $\mathbf{a}_3$, and along the way define some quantities $r_{j,k}$ that arise in the process:

$$
\begin{aligned}
\widehat{\mathbf{q}}_1 &:= \mathbf{a}_1 \\
\mathbf{q}_1 &:= \frac{1}{\|\widehat{\mathbf{q}}_1\|}\widehat{\mathbf{q}}_1 = \frac{1}{r_{1,1}}\widehat{\mathbf{q}}_1 \\
\\
\widehat{\mathbf{q}}_2 &:= \mathbf{a}_2 - \mathbf{q}_1\mathbf{q}_1^T\mathbf{a}_2 \\
&= \mathbf{a}_2 - r_{1,2}\mathbf{q}_1 \\
\mathbf{q}_2 &:= \frac{1}{\|\widehat{\mathbf{q}}_2\|}\widehat{\mathbf{q}}_2 = \frac{1}{r_{2,2}}\widehat{\mathbf{q}}_2 \\
\\
\widehat{\mathbf{q}}_3 &:= \mathbf{a}_3 - \mathbf{q}_1\mathbf{q}_1^T\mathbf{a}_3 - \mathbf{q}_2\mathbf{q}_2^T\mathbf{a}_3 \\
&= \mathbf{a}_3 - r_{1,3}\mathbf{q}_1 - r_{2,3}\mathbf{q}_2 \\
\mathbf{q}_3 &:= \frac{1}{\|\widehat{\mathbf{q}}_3\|}\widehat{\mathbf{q}}_3 = \frac{1}{r_{3,3}}\widehat{\mathbf{q}}_3.
\end{aligned}
$$

To summarize, we have defined

$$
r_{j,k} = \begin{cases} \mathbf{q}_j^T\mathbf{a}_k, & j < k; \\ \|\widehat{\mathbf{q}}_j\|, & j = k; \\ 0, & j > k. \end{cases}
$$

With this notation, three steps of the Gram–Schmidt process become:

$$
\begin{aligned}
r_{1,1}\mathbf{q}_1 &= \mathbf{a}_1 \\
r_{2,2}\mathbf{q}_2 &= \mathbf{a}_2 - r_{1,2}\mathbf{q}_1 \\
r_{3,3}\mathbf{q}_3 &= \mathbf{a}_3 - r_{1,3}\mathbf{q}_1 - r_{2,3}\mathbf{q}_2,
\end{aligned}
$$

or, collecting the $\mathbf{a}_j$ vectors on the left hand side:

$$
\begin{aligned}
\mathbf{a}_1 &= r_{1,1}\mathbf{q}_1 + 0\mathbf{q}_2 + 0\mathbf{q}_3 \\
\mathbf{a}_2 &= r_{1,2}\mathbf{q}_1 + r_{2,2}\mathbf{q}_2 + 0\mathbf{q}_3 \\
\mathbf{a}_3 &= r_{1,3}\mathbf{q}_1 + r_{2,3}\mathbf{q}_2 + r_{3,3}\mathbf{q}_3.
\end{aligned}
$$

Stack the orthonormal basis vectors as the columns of the *subunitary* matrix $\mathbf{Q} \in \mathbb{R}^{m\times n}$

$$
\mathbf{Q} = [\, \mathbf{q}_1 \quad \mathbf{q}_2 \quad \mathbf{q}_3 \,],
$$

and note that

$$
\mathbf{a}_1 = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 \end{bmatrix} \begin{bmatrix} r_{1,1} \\ 0 \\ 0 \end{bmatrix}
$$

$$
\mathbf{a}_2 = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 \end{bmatrix} \begin{bmatrix} r_{1,2} \\ r_{2,2} \\ 0 \end{bmatrix}
$$

$$\mathbf{a}_3 = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 \end{bmatrix} \begin{bmatrix} r_{1,3} \\ r_{2,3} \\ r_{3,3} \end{bmatrix},$$

which we organize in matrix form as

$$\begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \mathbf{a}_3 \end{bmatrix} = \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 \end{bmatrix} \begin{bmatrix} r_{1,1} & r_{1,2} & r_{1,3} \\ 0 & r_{2,2} & r_{2,3} \\ 0 & 0 & r_{3,3} \end{bmatrix}.$$

We summarize the entire process as:

$$\mathbf{A} = \mathbf{QR},$$

where $\mathbf{Q} \in \mathbb{R}^{m\times n}$ is subunitary and $\mathbf{R} \in \mathbb{R}^{n\times n}$ is upper triangular.

Now change your perspective: let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be any matrix with linearly independent columns, $m \geq n$. Performing the Gram–Schmidt process on those columns yields the decomposition $\mathbf{A} = \mathbf{QR}$, which is known as the *QR factorization*.

But wait, there's more! Suppose $\mathbf{b} \in \mathcal{R}(\mathbf{A})$, so there exists some $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{Ax} = \mathbf{b}$. Substitute the QR factorization to obtain $\mathbf{QRx} = \mathbf{b}$. Since we have written $\mathbf{b} = \mathbf{Q}(\mathbf{Rx})$, we see that $\mathbf{b} \in \mathcal{R}(\mathbf{Q})$: since $\mathbf{b}$ was any vector in the column space of $\mathbf{A}$, we conclude that $\mathcal{R}(\mathbf{A}) \subset \mathcal{R}(\mathbf{Q})$. Similarly, if $\mathbf{b} \in \mathcal{R}(\mathbf{Q})$, then there exists $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{Qx} = \mathbf{b}$. Since $\mathbf{R}$ is invertible, splice $\mathbf{I} = \mathbf{RR}^{-1}$ into this last equation to obtain $\mathbf{QR}(\mathbf{R}^{-1}\mathbf{x}) = \mathbf{b}$: so $\mathbf{A}(\mathbf{R}^{-1}\mathbf{x}) = \mathbf{b}$, hence $\mathbf{b} \in \mathcal{R}(\mathbf{A})$, and $\mathcal{R}(\mathbf{Q}) \subset \mathcal{R}(\mathbf{A})$. Since $\mathcal{R}(\mathbf{Q})$ and $\mathcal{R}(\mathbf{A})$ each contain the other, they must be equal.

Finally, notice that $\mathbf{P} = \mathbf{QQ}^T \in \mathbb{R}^{m\times m}$ is an orthogonal projector:

$$\mathbf{P}^2 = \mathbf{QQ}^T\mathbf{QQ}^T = \mathbf{QQ}^T = \mathbf{P},$$

since the columns of $\mathbf{Q}$ are orthonormal, $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$, and $(\mathbf{QQ}^T)^T = (\mathbf{Q}^T)^T\mathbf{Q}^T = \mathbf{QQ}^T$. In fact, one can confirm that

$$\mathbf{QQ}^T = \begin{bmatrix} \mathbf{q}_1 & \cdots & \mathbf{q}_n \end{bmatrix} \begin{bmatrix} \mathbf{q}_1^T \\ \vdots \\ \mathbf{q}_n^T \end{bmatrix} = \mathbf{q}_1\mathbf{q}_1^T + \cdots \mathbf{q}_n\mathbf{q}_n^T.$$

Notice that $\mathcal{R}(\mathbf{P}) = \mathcal{R}(\mathbf{Q})$, so $\mathbf{P}$ projects onto $\mathcal{R}(\mathbf{Q}) = \mathcal{R}(\mathbf{A})$.

Theorem 7 *Let* $\mathbf{A} = \mathbf{QR}$ *be a QR factorization of a matrix* $\mathbf{A} \in \mathbb{R}^{m\times n}$ *with linearly independent columns, with* $\mathbf{Q} \in \mathbb{R}^{m\times n}$ *and* $\mathbf{R} \in \mathbb{R}^{n\times n}$*. Then* $\mathcal{R}(\mathbf{Q}) = \mathcal{R}(\mathbf{A})$ *and* $\mathbf{QQ}^T$ *is an orthogonal projector onto* $\mathcal{R}(\mathbf{A})$*.*

5.4 QR solves systems

Suppose we wish to solve $\mathbf{Ax} = \mathbf{b}$ for a square matrix $\mathbf{A}$ with linearly independent columns. The linear independence implies that

$\mathbf{A}$ is invertible. Rather than perform Gaussian elimination, we could alternatively replace $\mathbf{A}$ with its QR factorization,

$$\mathbf{QRx} = \mathbf{b}.$$

Since $\mathbf{Q}$ is subunitary, $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$, so the $\mathbf{Q}$ term on the left can be cleared by premultiplying by $\mathbf{Q}^T$:

$$\mathbf{Rx} = \mathbf{Q}^T\mathbf{b}. \tag{5.1}$$

Since $\mathbf{R}$ is upper triangular, this $\mathbf{Rx} = \mathbf{Q}^T\mathbf{b}$ system can be solved for $\mathbf{x}$ by back-substitution, just as with the last step in Gaussian elimination. We can formally write the solution as

In general, the $\mathbf{R}$ factor in a QR factorization will be different from the $\mathbf{U}$ factor in an LU factorization, though both are upper triangular matrices.

$$\mathbf{x} = \mathbf{R}^{-1}\mathbf{Q}^T\mathbf{b}.$$

But is $\mathbf{R}$ invertible? Yes: since the columns of $\mathbf{A}$ are linearly independent, the diagonal entries of $\mathbf{R}$ obtained via the Gram–Schmidt process are always nonzero.

STUDENT EXPERIMENTS

5.23. The discussion above assumed $\mathbf{A}$ was a square matrix. Suppose $\mathbf{A} \in \mathbb{R}^{m\times n}$ with $m > n$, and $\mathbf{A}$ has n linearly independent columns. If $\mathbf{b} \in \mathcal{R}(\mathbf{A})$, write down a solution $\mathbf{x}$ to $\mathbf{Ax} = \mathbf{b}$ in terms of the QR factorization of $\mathbf{A}$. Is this solution unique?

5.5 *Least squares, take two*

Suppose that, instead of solving $\mathbf{Ax} = \mathbf{b}$, we instead wish to find $\mathbf{x}$ to minimize

$$\min_{\mathbf{x}\in\mathbb{C}^n} \|\mathbf{b} - \mathbf{Ax}\|,$$

the *least squares* problem discussed in the last chapter. We saw there that $\mathbf{x}$ minimized the misfit $\mathbf{b} - \mathbf{Ax}$ when $\mathbf{A}^T\mathbf{Ax} = \mathbf{A}^T\mathbf{b}$, and that this $\mathbf{x}$ was unique provided $\mathcal{N}(\mathbf{A}) = \{\mathbf{0}\}$.

Throughout this chapter we have assumed that $\mathbf{a}_1, \ldots, \mathbf{a}_n$ are linearly independent, giving the matrix

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \cdots & \mathbf{a}_n \end{bmatrix}$$

linearly independent columns: hence $\mathcal{N}(\mathbf{A}) = \{\mathbf{0}\}$, and the least squares solution is the unique $\mathbf{x}$ that satisfies

$$\mathbf{A}^T\mathbf{Ax} = \mathbf{A}^T\mathbf{b}.$$

Substitute $\mathbf{A} = \mathbf{QR}$ for $\mathbf{A}$ to get

$$(\mathbf{QR})^T(\mathbf{QR})\mathbf{x} = (\mathbf{QR})^T\mathbf{b}$$

which is equivalent to

$$\mathbf{R}^T\mathbf{Q}^T\mathbf{Q}\mathbf{R}\mathbf{x} = \mathbf{R}^T\mathbf{Q}^T\mathbf{b}.$$

Since $\mathbf{Q}$ is subunitary (its columns are orthogonal),

$$\mathbf{Q}^T\mathbf{Q} = \mathbf{I},$$

so our $\mathbf{x}$ solves

$$\mathbf{R}^T\mathbf{R}\mathbf{x} = \mathbf{R}^T\mathbf{Q}^T\mathbf{b}. \tag{5.2}$$

Note that $\mathbf{R}^T$ is a lower-triangular matrix with nonzero entries on the main diagonal, so it is invertible. Multiply both sides of (5.2) by $(\mathbf{R}^*)^{-1}$ to obtain

$$\mathbf{R}\mathbf{x} = \mathbf{Q}^T\mathbf{b}, \tag{5.3}$$

which, remarkably enough, is the same equation we obtained in (5.1) for solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ when $\mathbf{b} \in \mathcal{R}(\mathbf{A})$.

MATLAB captures this idea in its 'backslash' command, \. When you type A\b, you will obtain $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ when $\mathbf{A}$ is a square invertible matrix, and the $\mathbf{x}$ that minimizes $\|\mathbf{b} - \mathbf{A}\mathbf{x}\|$ when $\mathbf{A} \in \mathbb{R}^{m\times n}$ with $m > n$ having linearly independent columns.

But is the agreement of equations (5.1) and (5.3) really so remarkable? Insert $\mathbf{I} = \mathbf{Q}^T\mathbf{Q}$ on the right-hand side of (5.1) to obtain

$$\begin{aligned}\mathbf{R}\mathbf{x} &= \mathbf{Q}^T\mathbf{Q}\mathbf{Q}^T\mathbf{b}\\ &= \mathbf{Q}^T(\mathbf{Q}\mathbf{Q}^T)\mathbf{b}\\ &= \mathbf{Q}^T\mathbf{b}_R,\end{aligned}$$

noting that $\mathbf{Q}\mathbf{Q}^T$ is the orthogonal projector onto $\mathcal{R}(\mathbf{A})$, and $\mathbf{b}_R = \mathbf{Q}\mathbf{Q}^T\mathbf{b}$ is the vector given by the FTLA decomposition:

$$\begin{aligned}\mathbf{b} &= (\mathbf{Q}\mathbf{Q}^T\mathbf{b}) + (\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\mathbf{b}\\ &= \mathbf{b}_R + \mathbf{b}_N.\end{aligned}$$

So implicitly, equation (5.1) projects $\mathbf{b} \in \mathbb{R}^n$ to $\mathbf{b}_R$, to give an equation $\mathbf{R}\mathbf{x} = \mathbf{Q}^T\mathbf{b}_R$ that has a solution $\mathbf{x}$.