Example 3.4 Orthogonality and Projections. We illustrate these concepts by continuing with the least squares problem in Examples 3.1 and 3.3. At the solution $\boldsymbol{x}^T = [1236,\ 1943,\ 2416]$, the residual vector,

$$\boldsymbol{r} = \boldsymbol{b} - \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} - \boldsymbol{y} = \begin{bmatrix} 1237 \\ 1941 \\ 2417 \\ 711 \\ 1177 \\ 475 \end{bmatrix} - \begin{bmatrix} 1236 \\ 1943 \\ 2416 \\ 707 \\ 1180 \\ 473 \end{bmatrix} = \begin{bmatrix} 1 \\ -2 \\ 1 \\ 4 \\ -3 \\ 2 \end{bmatrix},$$

is orthogonal to each column of $\boldsymbol{A}$, i.e., $\boldsymbol{A}^T\boldsymbol{r} = \boldsymbol{o}$. The orthogonal projector onto $\text{span}(\boldsymbol{A})$ is given by

$$\boldsymbol{P} = \boldsymbol{A}(\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T = \frac{1}{4}\begin{bmatrix} 2 & 1 & 1 & -1 & -1 & 0 \\ 1 & 2 & 1 & 1 & 0 & -1 \\ 1 & 1 & 2 & 0 & 1 & 1 \\ -1 & 1 & 0 & 2 & 1 & -1 \\ -1 & 0 & 1 & 1 & 2 & 1 \\ 0 & -1 & 1 & -1 & 1 & 2 \end{bmatrix},$$

and the orthogonal projector onto $\text{span}(\boldsymbol{A})^{\perp}$ is given by

$$\boldsymbol{P}_{\perp} = \boldsymbol{I} - \boldsymbol{P} = \frac{1}{4}\begin{bmatrix} 2 & -1 & -1 & 1 & 1 & 0 \\ -1 & 2 & -1 & -1 & 0 & 1 \\ -1 & -1 & 2 & 0 & -1 & -1 \\ 1 & -1 & 0 & 2 & -1 & 1 \\ 1 & 0 & -1 & -1 & 2 & -1 \\ 0 & 1 & -1 & 1 & -1 & 2 \end{bmatrix},$$

so that $\boldsymbol{b} = \boldsymbol{P}\boldsymbol{b} + \boldsymbol{P}_{\perp}\boldsymbol{b} = \boldsymbol{y} + \boldsymbol{r}$.

3.3 Sensitivity and Conditioning

We turn now to the sensitivity and conditioning of linear least squares problems. First, we must extend the notion of matrix condition number to include rectangular matrices. The definition of condition number for a square matrix given in Section 2.3.3 makes use of the matrix inverse. A nonsquare matrix $\boldsymbol{A}$ does not have an inverse in the conventional sense, but it is possible to define a *pseudoinverse*, denoted by $\boldsymbol{A}^+$, that behaves like an inverse in many respects (see Exercise 3.32). We will later see a more general definition that applies to any matrix, but for now we consider only matrices $\boldsymbol{A}$ with full column rank, in which case $\boldsymbol{A}^T\boldsymbol{A}$ is nonsingular and we define the pseudoinverse of $\boldsymbol{A}$ to be

$$\boldsymbol{A}^+ = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T.$$

Trivially, we see that $\boldsymbol{A}^+\boldsymbol{A} = \boldsymbol{I}$, and from Section 3.2.2 we see that $\boldsymbol{P} = \boldsymbol{A}\boldsymbol{A}^+$ is an orthogonal projector onto span($\boldsymbol{A}$), so that the solution to the least squares problem $\boldsymbol{A}\boldsymbol{x} \cong \boldsymbol{b}$ is given by

$$\boldsymbol{x} = \boldsymbol{A}^+\boldsymbol{b}.$$

We now define the condition number of an $m \times n$ matrix with $\text{rank}(\boldsymbol{A}) = n$ to be

$$\text{cond}(\boldsymbol{A}) = \|\boldsymbol{A}\|_2 \cdot \|\boldsymbol{A}^+\|_2.$$

By convention, $\text{cond}(\boldsymbol{A}) = \infty$ if $\text{rank}(\boldsymbol{A}) < n$. Just as the condition number of a square matrix measures closeness to singularity, the condition number of a rectangular matrix measures closeness to rank deficiency.

Whereas the conditioning of a square linear system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ depends only on the matrix $\boldsymbol{A}$, the conditioning of a least squares problem $\boldsymbol{A}\boldsymbol{x} \cong \boldsymbol{b}$ depends on the right-hand-side vector $\boldsymbol{b}$ as well as the matrix $\boldsymbol{A}$, and thus $\text{cond}(\boldsymbol{A})$ alone does not suffice to characterize sensitivity. In particular, if $\boldsymbol{b}$ lies near span($\boldsymbol{A}$), then a small perturbation in $\boldsymbol{b}$ changes $\boldsymbol{y} = \boldsymbol{P}\boldsymbol{b}$ relatively little. But if $\boldsymbol{b}$ is nearly orthogonal to span($\boldsymbol{A}$), on the other hand, then $\boldsymbol{y} = \boldsymbol{P}\boldsymbol{b}$ itself will be relatively small, so that a small change in $\boldsymbol{b}$ can cause a relatively large change in $\boldsymbol{y}$, and hence in the least squares solution $\boldsymbol{x}$. Thus, for a given $\boldsymbol{A}$, we would expect a least squares problem with a $\boldsymbol{b}$ that yields a large residual (i.e., a poor fit to the data) to be more sensitive than one with a small residual (i.e., a good fit to the data). An appropriate measure of the closeness of $\boldsymbol{b}$ to span($\boldsymbol{A}$) is the ratio

$$\frac{\|\boldsymbol{A}\boldsymbol{x}\|_2}{\|\boldsymbol{b}\|_2} = \frac{\|\boldsymbol{y}\|_2}{\|\boldsymbol{b}\|_2} = \cos(\theta),$$

where θ is the angle between $\boldsymbol{b}$ and $\boldsymbol{y}$ (see Fig. 3.2). Thus, we expect greater sensitivity when this ratio is small, so that θ is near $\pi/2$.

We now make a more quantitative assessment of the sensitivity of the solution $\boldsymbol{x}$ of a least squares problem $\boldsymbol{A}\boldsymbol{x} \cong \boldsymbol{b}$, where $\boldsymbol{A}$ has full column rank. For simplicity, we will consider perturbations in $\boldsymbol{b}$ and $\boldsymbol{A}$ separately. For a perturbed right-hand-side vector $\boldsymbol{b} + \Delta\boldsymbol{b}$, the perturbed solution is given by the normal equations

$$\boldsymbol{A}^T\boldsymbol{A}(\boldsymbol{x} + \Delta\boldsymbol{x}) = \boldsymbol{A}^T(\boldsymbol{b} + \Delta\boldsymbol{b}).$$

Because $\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = \boldsymbol{A}^T\boldsymbol{b}$, we then have

$$\boldsymbol{A}^T\boldsymbol{A}\,\Delta\boldsymbol{x} = \boldsymbol{A}^T\Delta\boldsymbol{b},$$

so that

$$\Delta\boldsymbol{x} = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T\Delta\boldsymbol{b} = \boldsymbol{A}^+\Delta\boldsymbol{b}.$$

Taking norms, we obtain

$$\|\Delta\boldsymbol{x}\|_2 \leq \|\boldsymbol{A}^+\|_2 \cdot \|\Delta\boldsymbol{b}\|_2.$$

Dividing both sides by $\|\boldsymbol{x}\|_2$, we obtain the bound

$$\begin{aligned}
\frac{\|\Delta \boldsymbol{x}\|_2}{\|\boldsymbol{x}\|_2} &\leq \|\boldsymbol{A}^+\|_2 \frac{\|\Delta \boldsymbol{b}\|_2}{\|\boldsymbol{x}\|_2} \\
&= \operatorname{cond}(\boldsymbol{A}) \frac{\|\boldsymbol{b}\|_2}{\|\boldsymbol{A}\|_2 \cdot \|\boldsymbol{x}\|_2} \frac{\|\Delta \boldsymbol{b}\|_2}{\|\boldsymbol{b}\|_2} \\
&\leq \operatorname{cond}(\boldsymbol{A}) \frac{\|\boldsymbol{b}\|_2}{\|\boldsymbol{A}\boldsymbol{x}\|_2} \frac{\|\Delta \boldsymbol{b}\|_2}{\|\boldsymbol{b}\|_2} \\
&= \operatorname{cond}(\boldsymbol{A}) \frac{1}{\cos(\theta)} \frac{\|\Delta \boldsymbol{b}\|_2}{\|\boldsymbol{b}\|_2}.
\end{aligned}$$

Thus, the condition number for the least squares solution $\boldsymbol{x}$ with respect to perturbations in $\boldsymbol{b}$ depends on $\operatorname{cond}(\boldsymbol{A})$ and also on the angle θ between $\boldsymbol{b}$ and $\boldsymbol{A}\boldsymbol{x}$ (see Fig. 3.2). In particular, the condition number is approximately $\operatorname{cond}(\boldsymbol{A})$ when the residual is small, so that $\cos(\theta) \approx 1$, but the condition number can be arbitrarily worse than $\operatorname{cond}(\boldsymbol{A})$ when the residual is large, so that $\cos(\theta) \approx 0$.

For a perturbed matrix $\boldsymbol{A} + \boldsymbol{E}$, the perturbed solution is given by the normal equations

$$(\boldsymbol{A} + \boldsymbol{E})^T(\boldsymbol{A} + \boldsymbol{E})(\boldsymbol{x} + \Delta \boldsymbol{x}) = (\boldsymbol{A} + \boldsymbol{E})^T \boldsymbol{b}.$$

Noting that $\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = \boldsymbol{A}^T\boldsymbol{b}$, dropping second-order terms (i.e., products of small perturbations), and rearranging, we then have

$$\begin{aligned}
\boldsymbol{A}^T\boldsymbol{A}\,\Delta \boldsymbol{x} &\approx \boldsymbol{E}^T\boldsymbol{b} - \boldsymbol{E}^T\boldsymbol{A}\boldsymbol{x} - \boldsymbol{A}^T\boldsymbol{E}\boldsymbol{x} \\
&= \boldsymbol{E}^T(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}) - \boldsymbol{A}^T\boldsymbol{E}\boldsymbol{x} \\
&= \boldsymbol{E}^T\boldsymbol{r} - \boldsymbol{A}^T\boldsymbol{E}\boldsymbol{x},
\end{aligned}$$

so that

$$\Delta \boldsymbol{x} \approx (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{E}^T\boldsymbol{r} - (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T\boldsymbol{E}\boldsymbol{x} = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{E}^T\boldsymbol{r} - \boldsymbol{A}^+\boldsymbol{E}\boldsymbol{x}.$$

Taking norms, we obtain

$$\|\Delta \boldsymbol{x}\|_2 \lessapprox \|(\boldsymbol{A}^T\boldsymbol{A})^{-1}\|_2 \cdot \|\boldsymbol{E}\|_2 \cdot \|\boldsymbol{r}\|_2 + \|\boldsymbol{A}^+\|_2 \cdot \|\boldsymbol{E}\|_2 \cdot \|\boldsymbol{x}\|_2.$$

Dividing both sides by $\|\boldsymbol{x}\|_2$ and using the fact that $\|\boldsymbol{A}\|_2^2 \cdot \|(\boldsymbol{A}^T\boldsymbol{A})^{-1}\|_2 = [\operatorname{cond}(\boldsymbol{A})]^2$, we obtain the bound

$$\begin{aligned}
\frac{\|\Delta \boldsymbol{x}\|_2}{\|\boldsymbol{x}\|_2} &\lessapprox \|(\boldsymbol{A}^T\boldsymbol{A})^{-1}\|_2 \cdot \|\boldsymbol{E}\|_2 \frac{\|\boldsymbol{r}\|_2}{\|\boldsymbol{x}\|_2} + \|\boldsymbol{A}^+\|_2 \cdot \|\boldsymbol{E}\|_2 \\
&= [\operatorname{cond}(\boldsymbol{A})]^2 \frac{\|\boldsymbol{E}\|_2}{\|\boldsymbol{A}\|_2} \frac{\|\boldsymbol{r}\|_2}{\|\boldsymbol{A}\|_2 \cdot \|\boldsymbol{x}\|_2} + \operatorname{cond}(\boldsymbol{A}) \frac{\|\boldsymbol{E}\|_2}{\|\boldsymbol{A}\|_2} \\
&\leq \left([\operatorname{cond}(\boldsymbol{A})]^2 \frac{\|\boldsymbol{r}\|_2}{\|\boldsymbol{A}\boldsymbol{x}\|_2} + \operatorname{cond}(\boldsymbol{A}) \right) \frac{\|\boldsymbol{E}\|_2}{\|\boldsymbol{A}\|_2} \\
&= \left([\operatorname{cond}(\boldsymbol{A})]^2 \tan(\theta) + \operatorname{cond}(\boldsymbol{A}) \right) \frac{\|\boldsymbol{E}\|_2}{\|\boldsymbol{A}\|_2}.
\end{aligned}$$

Thus, the condition number for the least squares solution $\boldsymbol{x}$ with respect to perturbations in $\boldsymbol{A}$ depends on $\text{cond}(\boldsymbol{A})$ and also on the angle θ between $\boldsymbol{b}$ and $\boldsymbol{Ax}$ (see Fig. 3.2). In particular, the condition number is approximately $\text{cond}(\boldsymbol{A})$ when the residual is small, so that $\tan(\theta) \approx 0$, but the condition number is effectively squared for a moderate residual, and becomes arbitrarily large when the residual is larger still. These sensitivity results will not only enable us to assess the quality of least squares solutions, but will also play an important role in understanding the relative merits of the various algorithms for computing such solutions numerically.

Example 3.5 Sensitivity and Conditioning. We again illustrate these concepts by continuing with Examples 3.1, 3.3, and 3.4. The pseudoinverse is given by

$$\boldsymbol{A}^+ = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T = \frac{1}{4}\begin{bmatrix} 2 & 1 & 1 & -1 & -1 & 0 \\ 1 & 2 & 1 & 1 & 0 & -1 \\ 1 & 1 & 2 & 0 & 1 & 1 \end{bmatrix}.$$

The matrix norms can be computed to obtain

$$\|\boldsymbol{A}\|_2 = 2, \qquad \|\boldsymbol{A}^+\|_2 = 1,$$

so that

$$\text{cond}(\boldsymbol{A}) = \|\boldsymbol{A}\|_2 \cdot \|\boldsymbol{A}^+\|_2 = 2.$$

From the ratio

$$\cos(\theta) = \frac{\|\boldsymbol{Ax}\|_2}{\|\boldsymbol{b}\|_2} = \frac{\|\boldsymbol{y}\|_2}{\|\boldsymbol{b}\|_2} \approx \frac{3640.8761}{3640.8809} \approx 0.99999868,$$

we see that the angle θ between $\boldsymbol{b}$ and $\boldsymbol{y}$ is about 0.001625, which is very tiny, as expected for a problem with a very close fit to the data. From the small condition number and small angle θ, we conclude that this particular least squares problem is well-conditioned.

Example 3.6 Condition-Squaring Effect. Consider the matrix and perturbation

$$\boldsymbol{A} = \begin{bmatrix} 1 & 1 \\ \epsilon & -\epsilon \\ 0 & 0 \end{bmatrix}, \qquad \boldsymbol{E} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ -\epsilon & \epsilon \end{bmatrix},$$

where $\epsilon \ll 1$, say around $\sqrt{\epsilon_{\text{mach}}}$, for which we have

$$\text{cond}(\boldsymbol{A}) = 1/\epsilon, \qquad \|\boldsymbol{E}\|_2/\|\boldsymbol{A}\|_2 = \epsilon.$$

For the right-hand-side vector $\boldsymbol{b} = [1 \quad 0 \quad \epsilon]^T$, we have $\|\Delta\boldsymbol{x}\|_2/\|\boldsymbol{x}\|_2 = 0.5$, so the relative perturbation in the solution is about equal to $\text{cond}(\boldsymbol{A})$ times the relative perturbation in $\boldsymbol{A}$. There is no condition-squaring effect for this right-hand side because the residual is small and $\tan(\theta) \approx \epsilon$, effectively suppressing the condition-squared term in the perturbation bound.