

11.4 Least Squares and Minimum-Norm Solutions

Chapter 3 taught us to decide whether $\mathbf{Ax} = \mathbf{b}$ has a solution, and to find them all when it does. Real data rarely cooperates. Measure a quantity that depends linearly on two unknowns, take ten measurements, and you have ten equations in two unknowns; each measurement carries a little error, and the ten equations are almost certainly inconsistent. Refusing to answer is not useful. This section replaces the question “which $\mathbf{x}$ solves $\mathbf{Ax} = \mathbf{b}$?” with “which $\mathbf{x}$ comes closest?”, and shows that the new question always has an answer, computed by solving a square system.

The other half of the section handles the opposite complaint. When a system is consistent but underdetermined, there are too many solutions, and we want a canonical one. Orthogonality picks it: exactly one solution is perpendicular to the kernel, and it is the shortest.

Throughout, $F = \mathbb{R}$ or $F = \mathbb{C}$, and F^n and F^m carry the standard inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_i x_i \overline{y_i}$, with $\|\mathbf{x}\|^2 = \sum_i |x_i|^2$.

11.4.1 The residual, and a piece of notation

Given $\mathbf{A} \in M_{m \times n}(F)$ and $\mathbf{b} \in F^m$, write $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$ for the **residual** of a candidate $\mathbf{x}$. The system is consistent exactly when some $\mathbf{x}$ makes $\mathbf{r} = \mathbf{0}$. When none does, we ask for the $\mathbf{x}$ making $\mathbf{r}$ as short as possible, that is, minimizing

$$\|\mathbf{Ax} - \mathbf{b}\|^2 = \sum_{i=1}^m |(\mathbf{Ax})_i - b_i|^2,$$

a **sum of squares**; hence the name **least squares**.

Two observations turn this into geometry. First, as $\mathbf{x}$ ranges over F^n , the vector $\mathbf{Ax}$ ranges over exactly the column space $\text{col}(\mathbf{A}) \subseteq F^m$ (Theorem 0.9.3), a finite-dimensional subspace. Second, minimizing $\|\mathbf{Ax} - \mathbf{b}\|$ over all $\mathbf{x}$ is therefore the same as finding the point of $\text{col}(\mathbf{A})$ nearest to $\mathbf{b}$. We already know that point: it is $P_{\text{col}(\mathbf{A})}\mathbf{b}$, by Theorem 11.3.5. All that remains is to turn “the residual is perpendicular to $\text{col}(\mathbf{A})$ ” into equations, and for that we need one piece of notation.

Recall the conjugate transpose $\mathbf{A}^* \in M_{n \times m}(F)$ of Definition 11.1.2: the matrix with entries $(\mathbf{A}^*)_{ij} = \overline{a_{ji}}$, obtained by transposing and then conjugating every entry, so that $\mathbf{A}^* = \mathbf{A}^\top$ over $\mathbb{R}$. Comparing entries and using Theorem 0.9.5 together with the rules for conjugation (Theorem 0.6.2) gives

$$(\mathbf{A} + \mathbf{B})^* = \mathbf{A}^* + \mathbf{B}^*, \quad (c\mathbf{A})^* = \overline{c} \mathbf{A}^*, \quad (\mathbf{AB})^* = \mathbf{B}^* \mathbf{A}^*, \quad \mathbf{A}^{**} = \mathbf{A},$$

whenever the sizes match. The third is the one worth seeing done: $((\mathbf{AB})^*)_{ij} = \overline{(\mathbf{AB})_{ji}} = \overline{\sum_k a_{jk} b_{ki}} = \sum_k \overline{b_{ki}} \overline{a_{jk}} = \sum_k (\mathbf{B}^*)_{ik} (\mathbf{A}^*)_{kj} = (\mathbf{B}^* \mathbf{A}^*)_{ij}$, which is where the order reverses. The one consequence we really use is the following.

Lemma 11.4.1. Moving a matrix across the inner product

Let $\mathbf{A} \in M_{m \times n}(F)$. Then for all $\mathbf{x} \in F^n$ and $\mathbf{y} \in F^m$,

$$\langle \mathbf{Ax}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{A}^* \mathbf{y} \rangle, \quad \langle \mathbf{y}, \mathbf{Ax} \rangle = \langle \mathbf{A}^* \mathbf{y}, \mathbf{x} \rangle.$$

Proof. Regarding vectors as columns, $\langle \mathbf{u}, \mathbf{w} \rangle = \mathbf{w}^* \mathbf{u}$, a 1×1 matrix identified with its entry. Hence

$$\langle \mathbf{Ax}, \mathbf{y} \rangle = \mathbf{y}^* (\mathbf{Ax}) = (\mathbf{y}^* \mathbf{A}) \mathbf{x} = ((\mathbf{A}^* \mathbf{y})^*) \mathbf{x} = \langle \mathbf{x}, \mathbf{A}^* \mathbf{y} \rangle,$$

where the third equality uses $(\mathbf{A}^* \mathbf{y})^* = \mathbf{y}^* \mathbf{A}^{**} = \mathbf{y}^* \mathbf{A}$. The second identity follows by conjugating the first, since $\langle \mathbf{y}, \mathbf{Ax} \rangle = \overline{\langle \mathbf{Ax}, \mathbf{y} \rangle}$ and $\langle \mathbf{x}, \mathbf{A}^* \mathbf{y} \rangle = \overline{\langle \mathbf{A}^* \mathbf{y}, \mathbf{x} \rangle}$. ■

The lemma is the matrix case of a general principle: *to move a map across an inner product, pay with its adjoint*. Section 6 develops that principle for linear maps between abstract inner product spaces; here the two lines above are all we need.

11.4.2 The normal equations

Theorem 11.4.2. Least Squares Theorem

Let $\mathbf{A} \in M_{m \times n}(F)$ and $\mathbf{b} \in F^m$, and write $U = \text{col}(\mathbf{A})$. For $\mathbf{x} \in F^n$, the following are equivalent:

- (a) $\|\mathbf{Ax} - \mathbf{b}\| \leq \|\mathbf{Az} - \mathbf{b}\|$ for every $\mathbf{z} \in F^n$;
- (b) $\mathbf{Ax} = P_U \mathbf{b}$;
- (c) $\mathbf{A}^* \mathbf{Ax} = \mathbf{A}^* \mathbf{b}$.

Moreover such an $\mathbf{x}$ always exists. A vector satisfying these conditions is called a **least-squares solution** of $\mathbf{Ax} = \mathbf{b}$, and the equations in (c) are the **normal equations**.

Idea. Condition (a) is a minimization over $\mathbf{x}$; condition (b) is the same minimization rewritten in F^m , where Theorem 11.3.5 has already solved it. So (a) $\Leftrightarrow$ (b) is a translation, not an argument. The content is (b) $\Leftrightarrow$ (c): saying $\mathbf{Ax} = P_U \mathbf{b}$ is the same as saying the residual is perpendicular to U , and “perpendicular to every column of $\mathbf{A}$ ” is what $\mathbf{A}^*(\mathbf{b} - \mathbf{Ax}) = \mathbf{0}$ says. Existence is free, because $P_U \mathbf{b}$ lies in $U = \text{col}(\mathbf{A})$ and so is $\mathbf{Ax}$ for some $\mathbf{x}$.

Proof. Since $\{\mathbf{Az} : \mathbf{z} \in F^n\} = \text{col}(\mathbf{A}) = U$ (Theorem 0.9.3), condition (a) says that $\mathbf{Ax}$ is a vector of U at least as close to $\mathbf{b}$ as every other vector of U . By Theorem 11.3.5 the unique such vector is $P_U \mathbf{b}$. Hence (a) $\Leftrightarrow$ (b).

- (b) $\Leftrightarrow$ (c). Write $\mathbf{b} = P_U \mathbf{b} + (\mathbf{b} - P_U \mathbf{b})$, the splitting of Theorem 11.3.2 (a) with $\mathbf{b} - P_U \mathbf{b} \in U^\perp$. For $\mathbf{x} \in F^n$ we have $\mathbf{Ax} \in U$, so $\mathbf{b} - \mathbf{Ax} \in U^\perp$ holds if and only if $\mathbf{Ax}$ is the U -part of $\mathbf{b}$, that is, if and only if $\mathbf{Ax} = P_U \mathbf{b}$. It remains to see that

$$\mathbf{b} - \mathbf{Ax} \in U^\perp \iff \mathbf{A}^* \mathbf{Ax} = \mathbf{A}^* \mathbf{b}.$$

Put $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$. By Proposition 11.3.1, $\mathbf{r} \in U^\perp$ if and only if $\mathbf{r}$ is orthogonal to every $\mathbf{Az}$, and by Lemma 11.4.1,

$$\langle \mathbf{r}, \mathbf{Az} \rangle = \langle \mathbf{A}^* \mathbf{r}, \mathbf{z} \rangle \quad \text{for every } \mathbf{z} \in F^n.$$

If $\mathbf{A}^*\mathbf{r} = \mathbf{0}$, all these vanish. Conversely, if all vanish, take $\mathbf{z} = \mathbf{A}^*\mathbf{r}$ to get $\langle \mathbf{A}^*\mathbf{r}, \mathbf{A}^*\mathbf{r} \rangle = 0$, hence $\mathbf{A}^*\mathbf{r} = \mathbf{0}$ by positive definiteness. So $\mathbf{r} \in U^\perp$ if and only if $\mathbf{A}^*\mathbf{r} = \mathbf{0}$, which rearranges to $\mathbf{A}^*\mathbf{Ax} = \mathbf{A}^*\mathbf{b}$.

Existence. $P_U\mathbf{b} \in U = \text{col}(\mathbf{A})$, so $P_U\mathbf{b} = \mathbf{Ax}$ for some $\mathbf{x} \in F^n$, and that $\mathbf{x}$ satisfies (b). ■

Three things are worth extracting. The normal equations are a **square** $n \times n$ system, whatever the shape of $\mathbf{A}$. They are **always consistent**, by the existence clause, even when $\mathbf{Ax} = \mathbf{b}$ is not. And the minimizing *value* $\mathbf{Ax}$ is unique — it is $P_U\mathbf{b}$ — even if the minimizing $\mathbf{x}$ is not.

Warning. “Closest” is measured by the inner product, so the answer moves when the inner product does. Least squares is not a coordinate-free notion: it minimizes $\|\mathbf{Ax} - \mathbf{b}\|$ for a particular norm on F^m , and rescaling one coordinate — reporting one measurement in centimeters instead of meters, say — is a change of inner product. Concretely, for the data $(0, 2), (1, 1), (2, 8)$ the least-squares line is $y = \frac{2}{3} + 3t$; if the middle point is trusted twice as much, so that the second coordinate of F^3 is weighted by 2, the best line becomes $y = 3t$. Both are correct answers to their own question, and Exercise C1 works the general case out. Before fitting, decide which errors count equally.

Quick check. The system $\mathbf{Ax} = \mathbf{b}$ may have no solution at all, yet the normal equations $\mathbf{A}^*\mathbf{Ax} = \mathbf{A}^*\mathbf{b}$ always have one. Where does the proof arrange that, and what stops the same argument from solving $\mathbf{Ax} = \mathbf{b}$ itself?

Solution. The existence clause uses that $P_U\mathbf{b}$ lies in $U = \text{col}(\mathbf{A})$, and every vector of $\text{col}(\mathbf{A})$ is $\mathbf{Ax}$ for some $\mathbf{x}$. So the normal equations ask $\mathbf{Ax}$ to hit $P_U\mathbf{b}$, a target that is inside the column space by construction. The original system asks $\mathbf{Ax}$ to hit $\mathbf{b}$, which need not be in $\text{col}(\mathbf{A})$ at all. Projecting first is exactly the step that moves the target into reach; what is given up is the requirement $\mathbf{r} = \mathbf{0}$, which becomes $\mathbf{r} \perp \text{col}(\mathbf{A})$.

11.4.3 When is the least-squares solution unique?

The minimizers of $\|\mathbf{Ax} - \mathbf{b}\|$ are the solutions of the normal equations, so they form the solution set of a linear system: a coset $\mathbf{x}_0 + \text{nul}(\mathbf{A}^*\mathbf{A})$ (Theorem 3.1.3). To read off uniqueness we therefore need $\text{nul}(\mathbf{A}^*\mathbf{A})$, and the following small lemma computes it.

Lemma 11.4.3. The null space of $\mathbf{A}^*\mathbf{A}$

Let $\mathbf{A} \in M_{m \times n}(F)$. Then $\text{nul}(\mathbf{A}^*\mathbf{A}) = \text{nul}(\mathbf{A})$.

Proof. ($\supseteq$) If $\mathbf{Ax} = \mathbf{0}$, then $\mathbf{A}^*\mathbf{Ax} = \mathbf{A}^*\mathbf{0} = \mathbf{0}$.

($\subseteq$) Suppose $\mathbf{A}^*\mathbf{Ax} = \mathbf{0}$. Pair with $\mathbf{x}$ and move $\mathbf{A}^*$ across using Lemma 11.4.1:

$$0 = \langle \mathbf{A}^*(\mathbf{Ax}), \mathbf{x} \rangle = \langle \mathbf{Ax}, \mathbf{Ax} \rangle = \|\mathbf{Ax}\|^2.$$

Hence $\mathbf{A}\mathbf{x} = \mathbf{0}$ by positive definiteness. This proves the lemma. ■

This is the “pair it with itself” move once more, and it is the whole reason the normal equations behave well. It will return in Section 6 in the form $\ker(T^*T) = \ker T$ for a linear map T between inner product spaces, with the same two-line proof; the matrix statement above is the case we need now.

Corollary 11.4.4. Uniqueness of the least-squares solution

Let $\mathbf{A} \in M_{m \times n}(F)$ and $\mathbf{b} \in F^m$. The following are equivalent:

- (a) $\mathbf{A}\mathbf{x} = \mathbf{b}$ has exactly one least-squares solution (for one, equivalently for every, $\mathbf{b}$);
- (b) the columns of $\mathbf{A}$ are linearly independent, that is $\text{rank } \mathbf{A} = n$;
- (c) $\mathbf{A}^*\mathbf{A}$ is invertible.

In that case the least-squares solution is $\mathbf{x} = (\mathbf{A}^*\mathbf{A})^{-1}\mathbf{A}^*\mathbf{b}$.

Proof. By Theorem 11.4.2 the least-squares solutions are the solutions of $\mathbf{A}^*\mathbf{A}\mathbf{x} = \mathbf{A}^*\mathbf{b}$, a consistent system; by Theorem 3.1.3 its solution set is a coset of $\text{nul}(\mathbf{A}^*\mathbf{A})$, which is $\text{nul}(\mathbf{A})$ by Lemma 11.4.3. So the solution is unique, for one $\mathbf{b}$ or for all, exactly when $\text{nul}(\mathbf{A}) = \{\mathbf{0}\}$. Now $\text{nul}(\mathbf{A}) = \{\mathbf{0}\}$ says precisely that the only vanishing linear combination of the columns of $\mathbf{A}$ is the trivial one, that is, that the columns are independent; and by Theorem 3.5.7 this is the same as $\text{rank } \mathbf{A} = n$. This proves (a) $\Leftrightarrow$ (b).

- (b) $\Leftrightarrow$ (c). The matrix $\mathbf{A}^*\mathbf{A}$ is $n \times n$, so it is invertible if and only if $\text{nul}(\mathbf{A}^*\mathbf{A}) = \{\mathbf{0}\}$ (Theorem 3.4.4). By Lemma 11.4.3 this happens exactly when $\text{nul}(\mathbf{A}) = \{\mathbf{0}\}$.

Finally, if $\mathbf{A}^*\mathbf{A}$ is invertible, multiplying the normal equations on the left by $(\mathbf{A}^*\mathbf{A})^{-1}$ gives the stated formula. ■

Warning. $(\mathbf{A}^*\mathbf{A})^{-1}$ is not $\mathbf{A}^{-1}(\mathbf{A}^*)^{-1}$. The rule $(\mathbf{BC})^{-1} = \mathbf{C}^{-1}\mathbf{B}^{-1}$ needs $\mathbf{B}$ and $\mathbf{C}$ to be invertible, and in least squares $\mathbf{A}$ is usually not even square. Take

$$\mathbf{A} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad (\mathbf{A}^\top \mathbf{A})^{-1} = \frac{1}{3} \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}.$$

The left side exists; the symbols $\mathbf{A}^{-1}$ and $(\mathbf{A}^\top)^{-1}$ do not exist at all, since $\mathbf{A}$ is 3×2 . The formula $\mathbf{x} = (\mathbf{A}^*\mathbf{A})^{-1}\mathbf{A}^*\mathbf{b}$ must be read as one object $(\mathbf{A}^*\mathbf{A})^{-1}$ applied to $\mathbf{A}^*\mathbf{b}$, never simplified.

11.4.4 Fitting a line and a parabola

Example. The least-squares line through four points

Find the line $y = c + dt$ that minimizes the sum of the squared vertical errors at the four data points

$$(1, 3), \quad (2, 1), \quad (3, 5), \quad (4, 5).$$

Compute the residual and the minimum value of the sum of squares.

Solution. Asking for $c + dt_i = y_i$ at the four points is the system $\mathbf{Ax} = \mathbf{b}$ with

$$\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} c \\ d \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 3 \\ 1 \\ 5 \\ 5 \end{pmatrix},$$

and $\|\mathbf{Ax} - \mathbf{b}\|^2$ is exactly the sum of the squared vertical errors. The columns of $\mathbf{A}$ are independent (they are not multiples of each other), so [Corollary 11.4.4](#) gives a unique answer. Over $\mathbb{R}$, $\mathbf{A}^* = \mathbf{A}^\top$, and

$$\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 4 & 10 \\ 10 & 30 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{b} = \begin{pmatrix} 3 + 1 + 5 + 5 \\ 3 + 2 + 15 + 20 \end{pmatrix} = \begin{pmatrix} 14 \\ 40 \end{pmatrix}.$$

The normal equations $4c + 10d = 14$, $10c + 30d = 40$ simplify to $2c + 5d = 7$ and $c + 3d = 4$; subtracting twice the second from the first gives $-d = -1$, so $d = 1$ and $c = 1$. The least-squares line is

$$y = 1 + t.$$

Its values at $t = 1, 2, 3, 4$ are 2, 3, 4, 5, so the residual is

$$\mathbf{r} = \mathbf{b} - \mathbf{Ax} = (3, 1, 5, 5) - (2, 3, 4, 5) = (1, -2, 1, 0),$$

and the minimum sum of squares is $\|\mathbf{r}\|^2 = 1 + 4 + 1 + 0 = 6$.

Check. The residual must be orthogonal to both columns of $\mathbf{A}$: $1 - 2 + 1 + 0 = 0$ and $1 - 4 + 3 + 0 = 0$. Both hold, which confirms the normal equations.

The same machinery fits any linear combination of prescribed functions; only the columns of $\mathbf{A}$ change. “Linear” refers to the unknown coefficients, not to the shape of the curve.

Example. The least-squares parabola

Find the quadratic $y = a + bt + ct^2$ minimizing the sum of squared errors at

$$(-2, 7), \quad (-1, 6), \quad (0, 2), \quad (1, 0), \quad (2, 5).$$

Solution. Now $\mathbf{A}$ has columns $(1, 1, 1, 1, 1)$, (t_i) and (t_i^2) :

$$\mathbf{A} = \begin{pmatrix} 1 & -2 & 4 \\ 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 7 \\ 6 \\ 2 \\ 0 \\ 5 \end{pmatrix}.$$

Because the nodes are symmetric about 0, every odd power sums to zero: $\sum t_i = 0$ and $\sum t_i^3 = 0$. With $\sum t_i^2 = 10$ and $\sum t_i^4 = 34$,

$$\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 5 & 0 & 10 \\ 0 & 10 & 0 \\ 10 & 0 & 34 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{b} = \begin{pmatrix} 20 \\ -10 \\ 54 \end{pmatrix}.$$

The middle equation decouples: $10b = -10$, so $b = -1$. The outer two are $5a + 10c = 20$ and $10a + 34c = 54$; the first gives $a = 4 - 2c$, and substituting gives $40 + 14c = 54$, so $c = 1$ and $a = 2$. The least-squares parabola is

$$y = 2 - t + t^2.$$

Its values at the nodes are 8, 4, 2, 2, 4, so $\mathbf{r} = (-1, 2, 0, -2, 1)$ and the minimum sum of squares is $\|\mathbf{r}\|^2 = 1 + 4 + 0 + 4 + 1 = 10$.

Check. $\sum r_i = 0$, $\sum t_i r_i = 2 - 2 + 0 - 2 + 2 = 0$, and $\sum t_i^2 r_i = -4 + 2 + 0 - 2 + 4 = 0$: the residual is orthogonal to all three columns.

Two practical remarks. Choosing the nodes symmetrically, as above, made $\mathbf{A}^\top \mathbf{A}$ nearly diagonal and the arithmetic painless; the same trick appears again in Section 9 when the columns are trigonometric functions. And although the normal equations are the right way to *understand* least squares, they are not always the right way to *compute* it: Section 8 solves the same problem through a QR factorization, which avoids forming $\mathbf{A}^* \mathbf{A}$ at all.

11.4.5 The minimum-norm solution

Now the opposite situation: $\mathbf{Ax} = \mathbf{b}$ is consistent, but $\mathbf{A}$ has a nontrivial null space, so the solutions form an infinite family $\mathbf{x}_0 + \text{nul}(\mathbf{A})$ (Theorem 3.1.3). Which one should we report? Orthogonality answers.

Theorem 11.4.5. Minimum-Norm Solution

Let $\mathbf{A} \in M_{m \times n}(F)$ and let $\mathbf{b} \in F^m$ be such that $\mathbf{Ax} = \mathbf{b}$ is consistent. Write $K = \text{nul}(\mathbf{A})$. Then:

- (a) exactly one solution $\mathbf{x}_{\min}$ of $\mathbf{Ax} = \mathbf{b}$ lies in $K^\perp$;
- (b) $\|\mathbf{x}_{\min}\| < \|\mathbf{x}\|$ for every other solution $\mathbf{x}$.

We call $\mathbf{x}_{\min}$ the **minimum-norm solution**.

Idea. The solution set is a coset $\mathbf{x}_0 + K$, and $F^n = K \oplus K^\perp$. Throwing away the K -part of $\mathbf{x}_0$ keeps it a solution, because that part is in the kernel, and lands it in $K^\perp$. Once there, Pythagoras compares it with every other solution: the others differ from it by a kernel vector, which is perpendicular to it and only adds length.

Proof.

- (a) *Existence.* Let $\mathbf{x}_0$ be any solution. Since K is a subspace of the finite-dimensional space F^n , Theorem 11.3.2 (a) gives $F^n = K \oplus K^\perp$; write $\mathbf{x}_0 = \mathbf{k} + \mathbf{x}_{\min}$ with $\mathbf{k} \in K$ and $\mathbf{x}_{\min} \in K^\perp$. Then $\mathbf{A}\mathbf{x}_{\min} = \mathbf{A}\mathbf{x}_0 - \mathbf{A}\mathbf{k} = \mathbf{b} - \mathbf{0} = \mathbf{b}$, so $\mathbf{x}_{\min}$ is a solution lying in $K^\perp$.

Uniqueness. If $\mathbf{x}_{\min}$ and $\mathbf{y}$ are solutions in $K^\perp$, then $\mathbf{A}(\mathbf{x}_{\min} - \mathbf{y}) = \mathbf{0}$, so $\mathbf{x}_{\min} - \mathbf{y} \in K \cap K^\perp = \{\mathbf{0}\}$ by Theorem 11.3.2 (a).

- (b) Every solution has the form $\mathbf{x} = \mathbf{x}_{\min} + \mathbf{k}$ with $\mathbf{k} \in K$, by Theorem 3.1.3. Since $\mathbf{x}_{\min} \in K^\perp$ and $\mathbf{k} \in K$, the two are orthogonal, and Theorem 11.1.3 gives

$$\|\mathbf{x}\|^2 = \|\mathbf{x}_{\min}\|^2 + \|\mathbf{k}\|^2.$$

If $\mathbf{x} \neq \mathbf{x}_{\min}$, then $\mathbf{k} \neq \mathbf{0}$, so $\|\mathbf{k}\|^2 > 0$ and $\|\mathbf{x}\| > \|\mathbf{x}_{\min}\|$. This proves the theorem. ■

The subspace $K^\perp$ has a matrix description that makes $\mathbf{x}_{\min}$ computable. For $\mathbf{x} \in F^n$,

$$\begin{aligned} \mathbf{x} \in \text{col}(\mathbf{A}^*)^\perp &\iff \langle \mathbf{x}, \mathbf{A}^*\mathbf{y} \rangle = 0 \text{ for all } \mathbf{y} \in F^m \\ &\iff \langle \mathbf{A}\mathbf{x}, \mathbf{y} \rangle = 0 \text{ for all } \mathbf{y} \in F^m \\ &\iff \mathbf{A}\mathbf{x} = \mathbf{0}, \end{aligned}$$

where the first step is Definition 11.3.1 together with $\text{col}(\mathbf{A}^*) = \{\mathbf{A}^*\mathbf{y} : \mathbf{y} \in F^m\}$ (Theorem 0.9.3), the second is Lemma 11.4.1, and the last uses the choice $\mathbf{y} = \mathbf{A}\mathbf{x}$ with positive definiteness. So $\text{nul}(\mathbf{A}) = \text{col}(\mathbf{A}^*)^\perp$, and applying $\perp$ once more (Theorem 11.3.2 (b)) gives

$$(\text{nul}(\mathbf{A}))^\perp = \text{col}(\mathbf{A}^*).$$

In words: the minimum-norm solution is the unique solution that is a combination of the columns of $\mathbf{A}^*$, that is, of the conjugated rows of $\mathbf{A}$. Section 6 proves the same statement for linear maps, together with the three companion identities that make up the four fundamental subspaces.

Example. A minimum-norm solution

Find the solution of smallest norm of

$$\begin{cases} x_1 + x_2 + x_3 = 3, \\ x_1 + 2x_2 + 3x_3 = 7. \end{cases}$$

Solution. Here $\mathbf{A} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \end{pmatrix}$. Subtracting the first equation from the second gives $x_2 +$

$2x_3 = 4$, so $x_2 = 4 - 2x_3$ and then $x_1 = 3 - x_2 - x_3 = x_3 - 1$. With $t = x_3$ the solutions are

$$\mathbf{x}(t) = (-1, 4, 0) + t(1, -2, 1), \quad t \in \mathbb{R},$$

so $K = \text{nul}(\mathbf{A}) = \text{span}((1, -2, 1))$. By [Theorem 11.4.5](#) the answer is the unique t with $\mathbf{x}(t) \perp (1, -2, 1)$:

$$\langle \mathbf{x}(t), (1, -2, 1) \rangle = (-1 + t) - 2(4 - 2t) + t = 6t - 9 = 0, \quad t = \frac{3}{2}.$$

Hence $\mathbf{x}_{\min} = (\frac{1}{2}, 1, \frac{3}{2})$, with $\|\mathbf{x}_{\min}\|^2 = \frac{1}{4} + 1 + \frac{9}{4} = \frac{7}{2}$.

Check. It is a solution, since $\frac{1}{2} + 1 + \frac{3}{2} = 3$ and $\frac{1}{2} + 2 + \frac{9}{2} = 7$, and it is orthogonal to the kernel, since $\frac{1}{2} - 2 + \frac{3}{2} = 0$. By the identity $(\text{nul } \mathbf{A})^\perp = \text{col}(\mathbf{A}^\top)$ above it must also be a combination of the rows of $\mathbf{A}$, and it is: $\mathbf{x}_{\min} = \frac{1}{2}(1, 2, 3)$, half the second row. Finally, the solution $\mathbf{x}(0) = (-1, 4, 0)$ has $\|\mathbf{x}(0)\|^2 = 17 > \frac{7}{2}$, as [Theorem 11.4.5](#) (b) requires.

The two answers of this section — the least-squares solution and the minimum-norm solution — are usually wanted together: given any $\mathbf{A}$ and $\mathbf{b}$, first minimize $\|\mathbf{Ax} - \mathbf{b}\|$, then among the minimizers take the shortest. Chapter 13 packages the combined answer as a single matrix, the **pseudoinverse** $\mathbf{A}^+$, with $\mathbf{x} = \mathbf{A}^+\mathbf{b}$ in all cases at once.

11.4.6 Exercises

11.4.6.1 A. Check your understanding

Exercise. A1

- Define a least-squares solution of $\mathbf{Ax} = \mathbf{b}$, and write down the normal equations.
- True or false: if $\mathbf{Ax} = \mathbf{b}$ is consistent, then its solutions are exactly its least-squares solutions. Justify your answer.
- Explain in one sentence why the normal equations are always consistent.
- Give a condition on the columns of $\mathbf{A}$ equivalent to uniqueness of the least-squares solution.
- Where does the minimum-norm solution of a consistent system live, and why is it unique there?
- True or false: $(\mathbf{A}^*\mathbf{A})^{-1} = \mathbf{A}^{-1}(\mathbf{A}^*)^{-1}$. Justify your answer.

Solution.

- A vector $\mathbf{x} \in F^n$ with $\|\mathbf{Ax} - \mathbf{b}\| \leq \|\mathbf{Az} - \mathbf{b}\|$ for every $\mathbf{z} \in F^n$. The normal equations are $\mathbf{A}^*\mathbf{Ax} = \mathbf{A}^*\mathbf{b}$ ([Theorem 11.4.2](#)).
- True. If $\mathbf{Ax}_0 = \mathbf{b}$ then $\|\mathbf{Ax}_0 - \mathbf{b}\| = 0$, which is the smallest possible value, so $\mathbf{x}_0$ is a least-squares solution; conversely a least-squares solution $\mathbf{x}$ has $\|\mathbf{Ax} - \mathbf{b}\| \leq 0$, hence $\mathbf{Ax} = \mathbf{b}$. Least squares extends the notion of a solution without changing it when solutions exist.
- Because the target $P_{\text{col}(\mathbf{A})}\mathbf{b}$ lies in $\text{col}(\mathbf{A})$ by construction, so $\mathbf{Ax} = P_{\text{col}(\mathbf{A})}\mathbf{b}$ is solvable ([Theorem 11.4.2](#), existence).
- The columns are linearly independent, equivalently $\text{rank } \mathbf{A} = n$, equivalently $\mathbf{A}^*\mathbf{A}$ is invertible ([Corollary 11.4.4](#)).

- (e) In $(\text{nul } \mathbf{A})^\perp = \text{col}(\mathbf{A}^*)$. Two solutions there would differ by an element of $\text{nul}(\mathbf{A}) \cap (\text{nul } \mathbf{A})^\perp = \{\mathbf{0}\}$ (Theorem 11.4.5).
- (f) False, and usually meaningless: $\mathbf{A}$ need not be square, so $\mathbf{A}^{-1}$ need not exist, while $(\mathbf{A}^* \mathbf{A})^{-1}$ can exist perfectly well — see the 3×2 example in the warning after Corollary 11.4.4.

11.4.6.2 B. Practice

Exercise. B1: A least-squares line

Find the least-squares line $y = c + dt$ for the data $(0, 3), (1, 2), (2, 3), (3, 6)$. Compute the residual, verify that it is orthogonal to both columns of $\mathbf{A}$, and give the minimum sum of squares.

Solution. With $\mathbf{A}$ having columns $(1, 1, 1, 1)$ and $(0, 1, 2, 3)$ and $\mathbf{b} = (3, 2, 3, 6)$,

$$\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 4 & 6 \\ 6 & 14 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{b} = \begin{pmatrix} 14 \\ 0 + 2 + 6 + 18 \end{pmatrix} = \begin{pmatrix} 14 \\ 26 \end{pmatrix}.$$

The normal equations are $2c + 3d = 7$ and $3c + 7d = 13$. Multiplying the first by 3 and the second by 2 and subtracting gives $9d - 14d = 21 - 26$, so $d = 1$, and then $c = 2$. The line is $y = 2 + t$.

Its values at $t = 0, 1, 2, 3$ are $2, 3, 4, 5$, so $\mathbf{r} = (1, -1, -1, 1)$. Orthogonality: $1 - 1 - 1 + 1 = 0$ against the first column and $0 - 1 - 2 + 3 = 0$ against the second. The minimum sum of squares is $\|\mathbf{r}\|^2 = 4$.

Exercise. B2: A least-squares parabola

Find the least-squares quadratic $y = a + bt + ct^2$ for the data $(-1, 2), (0, 8), (1, 6), (2, 16)$, and the minimum sum of squares. *Hint: write down $\mathbf{A}^\top \mathbf{A}$ from the power sums $\sum t_i^k$, $k = 0, \dots, 4$.*

Solution. The power sums are $\sum 1 = 4$, $\sum t_i = 2$, $\sum t_i^2 = 6$, $\sum t_i^3 = 8$, $\sum t_i^4 = 18$, so

$$\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 4 & 2 & 6 \\ 2 & 6 & 8 \\ 6 & 8 & 18 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{b} = \begin{pmatrix} 2 + 8 + 6 + 16 \\ -2 + 0 + 6 + 32 \\ 2 + 0 + 6 + 64 \end{pmatrix} = \begin{pmatrix} 32 \\ 36 \\ 72 \end{pmatrix}.$$

Solving: from the first two equations, $2a + b + 3c = 16$ and $a + 3b + 4c = 18$; eliminating a gives $5b + 5c = 20$, so $b + c = 4$. The third equation is $3a + 4b + 9c = 36$. Substituting $b = 4 - c$ into $2a + b + 3c = 16$ gives $2a + 2c = 12$, so $a = 6 - c$; then $3(6 - c) + 4(4 - c) + 9c = 36$ becomes $34 + 2c = 36$, so $c = 1$, $b = 3$, $a = 5$. The parabola is

$$y = 5 + 3t + t^2.$$

Its values at $t = -1, 0, 1, 2$ are $3, 5, 9, 15$, so $\mathbf{r} = (-1, 3, -3, 1)$ and the minimum sum of

squares is $1 + 9 + 9 + 1 = 20$. Check: $\sum r_i = 0$, $\sum t_i r_i = 1 + 0 - 3 + 2 = 0$, $\sum t_i^2 r_i = -1 + 0 - 3 + 4 = 0$.

Exercise. B3: A minimum-norm solution

Find the minimum-norm solution of

$$\begin{cases} x_1 + x_3 = 2, \\ x_2 + x_3 = 3, \end{cases}$$

and verify that it is orthogonal to $\text{nul}(\mathbf{A})$.

Solution. With $x_3 = t$ the solutions are $\mathbf{x}(t) = (2 - t, 3 - t, t) = (2, 3, 0) + t(-1, -1, 1)$, so $\text{nul}(\mathbf{A}) = \text{span}((-1, -1, 1))$. By [Theorem 11.4.5](#) we need $\langle \mathbf{x}(t), (-1, -1, 1) \rangle = 0$:

$$-(2 - t) - (3 - t) + t = 3t - 5 = 0, \quad t = \frac{5}{3}.$$

So $\mathbf{x}_{\min} = (\frac{1}{3}, \frac{4}{3}, \frac{5}{3})$, with $\|\mathbf{x}_{\min}\|^2 = \frac{1+16+25}{9} = \frac{14}{3}$. It solves the system, since $\frac{1}{3} + \frac{5}{3} = 2$ and $\frac{4}{3} + \frac{5}{3} = 3$, and $-\frac{1}{3} - \frac{4}{3} + \frac{5}{3} = 0$ confirms the orthogonality. For comparison, $\|(2, 3, 0)\|^2 = 13 > \frac{14}{3}$.

11.4.6.3 C. Going deeper

Exercise. C1: Weighted least squares

Let $\mathbf{A} \in M_{m \times n}(\mathbb{R})$, $\mathbf{b} \in \mathbb{R}^m$, and let $w_1, \dots, w_m > 0$ be weights, with $\mathbf{W} = \text{diag}(w_1, \dots, w_m)$.

- Show that $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbf{W}} = \mathbf{y}^\top \mathbf{W} \mathbf{x}$ is an inner product on $\mathbb{R}^m$.
- Prove that the minimizers of $\sum_{i=1}^m w_i ((\mathbf{A} \mathbf{x})_i - b_i)^2$ are exactly the solutions of $\mathbf{A}^\top \mathbf{W} \mathbf{A} \mathbf{x} = \mathbf{A}^\top \mathbf{W} \mathbf{b}$.
- For the data $(0, 2), (1, 1), (2, 8)$ with weights $(1, 2, 1)$, find the weighted least-squares line, and compare it with the unweighted one.

Solution.

- It is bilinear and symmetric because $\mathbf{y}^\top \mathbf{W} \mathbf{x} = \sum_i w_i x_i y_i$, and $\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{W}} = \sum_i w_i x_i^2 > 0$ whenever some $x_i \neq 0$, because every $w_i > 0$. (This is where positivity of the weights is used; a zero weight would destroy definiteness.)
- Write $\|\cdot\|_{\mathbf{W}}$ for the norm induced by $\langle \cdot, \cdot \rangle_{\mathbf{W}}$. The quantity to minimize is exactly $\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_{\mathbf{W}}^2$. The proof of [Theorem 11.4.2](#) used nothing about the inner product except its axioms, so it applies verbatim with $\langle \cdot, \cdot \rangle_{\mathbf{W}}$ in place of the standard one: the minimizers are the $\mathbf{x}$ for which $\mathbf{A} \mathbf{x} - \mathbf{b}$ is $\langle \cdot, \cdot \rangle_{\mathbf{W}}$ -orthogonal to $\text{col}(\mathbf{A})$, that is, for which $\langle \mathbf{A} \mathbf{x} - \mathbf{b}, \mathbf{A} \mathbf{z} \rangle_{\mathbf{W}} = 0$ for all $\mathbf{z} \in \mathbb{R}^n$. Now

$$\langle \mathbf{A} \mathbf{x} - \mathbf{b}, \mathbf{A} \mathbf{z} \rangle_{\mathbf{W}} = (\mathbf{A} \mathbf{z})^\top \mathbf{W} (\mathbf{A} \mathbf{x} - \mathbf{b}) = \mathbf{z}^\top (\mathbf{A}^\top \mathbf{W} \mathbf{A} \mathbf{x} - \mathbf{A}^\top \mathbf{W} \mathbf{b}),$$

and this vanishes for all $\mathbf{z} \in \mathbb{R}^n$ exactly when $\mathbf{A}^\top \mathbf{W} \mathbf{A} \mathbf{x} - \mathbf{A}^\top \mathbf{W} \mathbf{b} = \mathbf{0}$, by taking $\mathbf{z}$ equal to that vector.

- (c) With $\mathbf{A}$ having columns $(1, 1, 1)$ and $(0, 1, 2)$, $\mathbf{b} = (2, 1, 8)$ and $\mathbf{W} = \text{diag}(1, 2, 1)$,

$$\mathbf{A}^\top \mathbf{W} \mathbf{A} = \begin{pmatrix} 4 & 4 \\ 4 & 6 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{W} \mathbf{b} = \begin{pmatrix} 2 + 2 + 8 \\ 0 + 2 + 16 \end{pmatrix} = \begin{pmatrix} 12 \\ 18 \end{pmatrix},$$

so $4c + 4d = 12$ and $4c + 6d = 18$, giving $d = 3$ and $c = 0$: the weighted line is $y = 3t$. Unweighted, $\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix}$ and $\mathbf{A}^\top \mathbf{b} = (11, 17)$, giving $3c + 3d = 11$, $3c + 5d = 17$, so $d = 3$ and $c = \frac{2}{3}$: the unweighted line is $y = \frac{2}{3} + 3t$. The middle point is trusted twice as much in the weighted fit, and the line moves down to pass closer to it. The weighted fit has weighted error $4 + 2 \cdot 4 + 4 = 16$, smaller than the weighted error $\frac{160}{9}$ of the unweighted line, while the unweighted fit has plain error $\frac{32}{3}$, smaller than the plain error 12 of the weighted line: each line wins the contest it was designed for.

Exercise. C2: The residual, two ways

Let $\mathbf{x}$ be a least-squares solution of $\mathbf{Ax} = \mathbf{b}$ and let $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$.

- Prove that $\|\mathbf{r}\|^2 = \|\mathbf{b}\|^2 - \|\mathbf{Ax}\|^2$, and also that $\|\mathbf{r}\|^2 = \langle \mathbf{b}, \mathbf{r} \rangle = \|\mathbf{b}\|^2 - \langle \mathbf{A}^* \mathbf{b}, \mathbf{x} \rangle$.
- Use the second formula to recompute the minimum sum of squares in [Example \(The least-squares line through four points\)](#) without finding the residual vector.

Solution.

- By [Theorem 11.4.2](#), $\mathbf{Ax} = P_U \mathbf{b}$ with $U = \text{col}(\mathbf{A})$, so $\mathbf{r} = \mathbf{b} - P_U \mathbf{b} \in U^\perp$ while $\mathbf{Ax} \in U$. The two are orthogonal, and $\mathbf{b} = \mathbf{Ax} + \mathbf{r}$, so Theorem 11.1.3 gives $\|\mathbf{b}\|^2 = \|\mathbf{Ax}\|^2 + \|\mathbf{r}\|^2$, which is the first identity. For the second, $\langle \mathbf{b}, \mathbf{r} \rangle = \langle \mathbf{Ax} + \mathbf{r}, \mathbf{r} \rangle = \langle \mathbf{Ax}, \mathbf{r} \rangle + \|\mathbf{r}\|^2 = \|\mathbf{r}\|^2$, since $\langle \mathbf{Ax}, \mathbf{r} \rangle = \overline{\langle \mathbf{r}, \mathbf{Ax} \rangle} = 0$ by orthogonality. Expanding the other way, and using [Lemma 11.4.1](#),

$$\langle \mathbf{b}, \mathbf{r} \rangle = \langle \mathbf{b}, \mathbf{b} - \mathbf{Ax} \rangle = \|\mathbf{b}\|^2 - \langle \mathbf{b}, \mathbf{Ax} \rangle = \|\mathbf{b}\|^2 - \langle \mathbf{A}^* \mathbf{b}, \mathbf{x} \rangle.$$

- In [Example \(The least-squares line through four points\)](#), $\|\mathbf{b}\|^2 = 9 + 1 + 25 + 25 = 60$, $\mathbf{A}^\top \mathbf{b} = (14, 40)$ and $\mathbf{x} = (1, 1)$, so

$$\|\mathbf{r}\|^2 = 60 - \langle (14, 40), (1, 1) \rangle = 60 - 54 = 6,$$

matching the value found there. Only the data already assembled for the normal equations is needed.

Exercise. C3: A rank-deficient least-squares problem

Let

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 2 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix}.$$

- Show that $\mathbf{Ax} = \mathbf{b}$ is inconsistent and that $\text{rank } \mathbf{A} = 2$.
- Describe **all** least-squares solutions.

- (c) Find the one of smallest norm, and the minimum sum of squares.

Solution.

- (a) The third column is the sum of the first two, so the columns are dependent and $\text{rank } \mathbf{A} \leq 2$; the first two columns are independent, so $\text{rank } \mathbf{A} = 2$ and $\text{col}(\mathbf{A}) = \text{span}((1, 0, 1), (0, 1, 1))$. A vector (b_1, b_2, b_3) lies in $\text{col}(\mathbf{A})$ exactly when $b_3 = b_1 + b_2$; here $0 \neq 1 + 2$, so $\mathbf{b} \notin \text{col}(\mathbf{A})$ and the system is inconsistent.
- (b) By [Theorem 11.4.2](#) the least-squares solutions solve $\mathbf{A}^\top \mathbf{A} \mathbf{x} = \mathbf{A}^\top \mathbf{b}$. Here

$$\mathbf{A}^\top \mathbf{A} = \begin{pmatrix} 2 & 1 & 3 \\ 1 & 2 & 3 \\ 3 & 3 & 6 \end{pmatrix}, \quad \mathbf{A}^\top \mathbf{b} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}.$$

Row reducing, the third row is the sum of the first two, and subtracting suitable multiples gives $x_1 + x_3 = 0$ and $x_2 + x_3 = 1$. Hence

$$\mathbf{x}(t) = (0, 1, 0) + t(-1, -1, 1), \quad t \in \mathbb{R},$$

a coset of $\text{nul}(\mathbf{A}) = \text{span}((-1, -1, 1)) = \text{nul}(\mathbf{A}^\top \mathbf{A})$, as [Lemma 11.4.3](#) predicts.

- (c) The minimizers form a coset of $\text{nul}(\mathbf{A})$, so the argument of [Theorem 11.4.5](#) applies to it verbatim: the shortest is the one orthogonal to $(-1, -1, 1)$. From $-(-t) - (1 - t) + t = 3t - 1 = 0$ we get $t = \frac{1}{3}$ and

$$\mathbf{x}_{\min} = \left(-\frac{1}{3}, \frac{2}{3}, \frac{1}{3}\right), \quad \|\mathbf{x}_{\min}\|^2 = \frac{1+4+1}{9} = \frac{2}{3}.$$

Every minimizer gives the same $\mathbf{A} \mathbf{x} = (0, 1, 1) = P_{\text{col}(\mathbf{A})} \mathbf{b}$, with residual $\mathbf{r} = (1, 1, -1)$ and minimum sum of squares $\|\mathbf{r}\|^2 = 3$. *Check:* $\mathbf{r}$ is orthogonal to $(1, 0, 1)$ and to $(0, 1, 1)$, since $1 - 1 = 0$ and $1 - 1 = 0$.