INTRODUCTION

TO

LINEAR

ALGEBRA

Fifth Edition

MANUAL FOR INSTRUCTORS

Gilbert Strang

Massachusetts Institute of Technology

math.mit.edu/linearalgebra

web.mit.edu/18.06

video lectures: ocw.mit.edu

math.mit.edu/∼gs

www.wellesleycambridge.com

email: linearalgebrabook@gmail.com

Wellesley - Cambridge Press

Box 812060

Wellesley, Massachusetts 02482

Problem Set 11.1, page 516

1 Without exchange, pivots .001 and 1000; with exchange, 1 and -1. When the pivot is larger than the entries below it, all $|\ell_{ij}| = \dfrac{|\text{entry}|}{|\text{pivot}|} \le 1$. $A = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & -1 \\ -1 & 1 & 1 \end{bmatrix}$.

2 The exact inverse of hilb(3) is $A^{-1} = \begin{bmatrix} 9 & -36 & 30 \\ -36 & 192 & -180 \\ 30 & -180 & 180 \end{bmatrix}$.

3 $A\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 11/6 \\ 13/12 \\ 47/60 \end{bmatrix} = \begin{bmatrix} 1.833 \\ 1.083 \\ 0.783 \end{bmatrix}$ compares with $A\begin{bmatrix} 0 \\ 6 \\ -3.6 \end{bmatrix} = \begin{bmatrix} 1.80 \\ 1.10 \\ 0.78 \end{bmatrix}$. $\|\Delta \boldsymbol{b}\| < .04$ but $\|\Delta \boldsymbol{x}\| > 6$.
The difference $(1,1,1) - (0,6,-3.6)$ is in a direction $\Delta \boldsymbol{x}$ that has $A\Delta \boldsymbol{x}$ near zero.

4 The largest $\|\boldsymbol{x}\| = \|A^{-1}\boldsymbol{b}\|$ is $\|A^{-1}\| = 1/\lambda_{\min}$ since $A^{\mathrm{T}} = A$; largest error $10^{-16}/\lambda_{\min}$.

5 Each row of U has at most w entries. Use w multiplications to substitute components of $\boldsymbol{x}$ (already known from below) and divide by the pivot. Total for n rows $< wn$.

6 The triangular L^{-1}, U^{-1}, R^{-1} need $\frac{1}{2}n^2$ multiplications. Q needs n^2 to multiply the right side by $Q^{-1} = Q^{\mathrm{T}}$. So $QR\boldsymbol{x} = \boldsymbol{b}$ takes 1.5 times longer than $LU\boldsymbol{x} = \boldsymbol{b}$.

7 $UU^{-1} = I$: Back substitution needs $\frac{1}{2}j^2$ multiplications on column j, using the j by j upper left block. Then $\frac{1}{2}(1^2 + 2^2 + \cdots + n^2) \approx \frac{1}{2}(\frac{1}{3}n^3) = $ total to find U^{-1}.

8 $\begin{bmatrix} 1 & 0 \\ 2 & 2 \end{bmatrix} \to \begin{bmatrix} 2 & 2 \\ 1 & 0 \end{bmatrix} \to \begin{bmatrix} 2 & 2 \\ 0 & -1 \end{bmatrix} = U$ with $P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ and $L = \begin{bmatrix} 1 & 0 \\ .5 & 1 \end{bmatrix}$;

$A \to \begin{bmatrix} 2 & 2 & 0 \\ 1 & 0 & 1 \\ 0 & 2 & 0 \end{bmatrix} \to \begin{bmatrix} 2 & 2 & 0 \\ 0 & -1 & 1 \\ 0 & 2 & 0 \end{bmatrix} \to \begin{bmatrix} 2 & 2 & 0 \\ 0 & 2 & 0 \\ 0 & -1 & 1 \end{bmatrix} \to \begin{bmatrix} 2 & 2 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} = U$ with

$P = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$ and $L = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ .5 & -.5 & 1 \end{bmatrix}$.

9 $A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}$ has cofactors $C_{13} = C_{31} = C_{24} = C_{42} = 1$ and $C_{14} = C_{41} = -1$. A^{-1} is a full matrix!

10 With 16-digit floating point arithmetic the errors $\|\boldsymbol{x} - \boldsymbol{x}_{\text{computed}}\|$ for $\varepsilon = 10^{-3}$, 10^{-6}, 10^{-9}, 10^{-12}, 10^{-15} are of order 10^{-16}, 10^{-11}, 10^{-7}, 10^{-4}, 10^{-3}.

11 (a) $\cos\theta = 1/\sqrt{10}$, $\sin\theta = -3/\sqrt{10}$, $R = \frac{1}{\sqrt{10}} \begin{bmatrix} 1 & 3 \\ -3 & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ 3 & 5 \end{bmatrix} = \frac{1}{\sqrt{10}} \begin{bmatrix} 10 & 14 \\ 0 & 8 \end{bmatrix}$.

(b) A has eigenvalues 4 and 2. Put one of the unit eigenvectors in row 1 of Q: either

$Q = \dfrac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $QAQ^{-1} = \begin{bmatrix} 2 & -4 \\ 0 & 4 \end{bmatrix}$ **or**

$Q = \dfrac{1}{\sqrt{10}} \begin{bmatrix} 1 & -3 \\ 3 & 1 \end{bmatrix}$ and $QAQ^{-1} = \begin{bmatrix} 4 & -4 \\ 0 & 2 \end{bmatrix}$.

12 When A is multiplied by a plane rotation Q_{ij}, this changes the $2n$ (not n^2) entries in rows i and j. Then multiplying on the right by $(Q_{ij})^{-1} = (Q_{ij})^{\mathrm{T}}$ changes the $2n$ entries in columns i and j.

13 $Q_{ij}A$ uses $4n$ multiplications (2 for each entry in rows i and j). By factoring out $\cos\theta$, the entries 1 and $\pm\tan\theta$ need only $2n$ multiplications, which leads to $\frac{2}{3}n^3$ for QR.

14 The $(2,1)$ entry of $Q_{21}A$ is $\frac{1}{3}(-\sin\theta + 2\cos\theta)$. This is zero if $\sin\theta = 2\cos\theta$ or $\tan\theta = 2$. Then the $2, 1, \sqrt{5}$ right triangle has $\sin\theta = 2/\sqrt{5}$ and $\cos\theta = 1/\sqrt{5}$.

Every 3 by 3 rotation with $\det Q = +1$ is the product of 3 plane rotations.

15 This problem shows how elimination is more expensive (the nonzero multipliers in L and LL are counted by **nnz**(L) and **nnz**(LL)) when we spoil the tridiagonal K by a random permutation.

If on the other hand we start with a poorly ordered matrix K, an improved ordering is found by the code **symamd** discussed in this section.

16 The "red-black ordering" puts rows and columns 1 to 10 in the odd-even order $1, 3, 5, 7, 9, 2, 4, 6, 8, 10$. When K is the $-1, 2, -1$ tridiagonal matrix, odd points are connected only to even points (and 2 stays on the diagonal, connecting every point to itself):

$$K = \begin{bmatrix} 2 & -1 & & \\ -1 & 2 & -1 & \\ & \cdot & \cdot & \cdot \\ & & -1 & 2 \end{bmatrix} \text{ and } PKP^{\mathrm{T}} = \begin{bmatrix} 2I & D \\ D^{\mathrm{T}} & 2I \end{bmatrix} \text{ with}$$

$$D = \begin{bmatrix} -1 & & & & \\ -1 & -1 & & & \\ 0 & -1 & -1 & & \\ & & -1 & -1 & \\ & & & -1 & -1 \end{bmatrix} \begin{matrix} 1 \text{ to } 2 \\ 3 \text{ to } 2, 4 \\ 5 \text{ to } 4, 6 \\ 7 \text{ to } 6, 8 \\ 9 \text{ to } 8, 10 \end{matrix}$$

17 Jeff Stuart's **Shake a Stick** activity has long sticks representing the graphs of two linear equations in the x-y plane. The matrix is nearly singular and Section 9.2 shows how to compute its condition number $c = \|A\|\|A^{-1}\| = \sigma_{\max}/\sigma_{\min} \approx 80,000$:

$$A = \begin{bmatrix} 1 & 1.0001 \\ 1 & 1.0000 \end{bmatrix} \quad \|A\| \approx 2 \quad A^{-1} = 10000 \begin{bmatrix} -1 & 1.0001 \\ 1 & -1 \end{bmatrix} \quad \begin{matrix} \|A^{-1}\| \approx 20000 \\ c \approx 40000. \end{matrix}$$

Problem Set 11.2, page 522

1 $\|A\| = 2$, $\|A^{-1}\| = 2$, $c = 4$; $\|A\| = 3$, $\|A^{-1}\| = 1$, $c = 3$; $\|A\| = 2 + \sqrt{2} = \lambda_{\max}$ for positive definite A, $\|A^{-1}\| = 1/\lambda_{\min}$, comd $= (2+\sqrt{2})/(2-\sqrt{2}) = 5.83$.

2 $\|A\| = 2$, $c = 1$; $\|A\| = \sqrt{2}$, $c = \infty$ (singular matrix); $A^{\mathrm{T}}A = 2I$, $\|A\| = \sqrt{2}$, $c = 1$.

3 For the first inequality replace $\boldsymbol{x}$ by $B\boldsymbol{x}$ in $\|A\boldsymbol{x}\| \leq \|A\|\|\boldsymbol{x}\|$; the second inequality is just $\|B\boldsymbol{x}\| \leq \|B\|\|\boldsymbol{x}\|$. Then $\|AB\| = \max(\|AB\boldsymbol{x}\|/\|\boldsymbol{x}\|) \leq \|A\|\|B\|$.

4 $1 = \|I\| = \|AA^{-1}\| \leq \|A\|\|A^{-1}\| = c(A)$.

5 If $\Lambda_{\max} = \Lambda_{\min} = 1$ then all $\Lambda_i = 1$ and $A = SIS^{-1} = I$. The only matrices with $\|A\| = \|A^{-1}\| = 1$ are *orthogonal matrices.*

6 All orthogonal matrices have norm 1, so $\|A\| \le \|Q\|\|R\| = \|R\|$ and in reverse $\|R\| \le \|Q^{-1}\|\|A\| = \|A\|$. Then $\|A\| = \|R\|$. Inequality is usual in $\|A\| < \|L\|\|U\|$ when $A^{\mathrm{T}}A \ne AA^{\mathrm{T}}$. Use **norm** on a random A.

7 The triangle inequality gives $\|A\boldsymbol{x} + B\boldsymbol{x}\| \le \|A\boldsymbol{x}\| + \|B\boldsymbol{x}\|$. Divide by $\|\boldsymbol{x}\|$ and take the maximum over all nonzero vectors to find $\|A + B\| \le \|A\| + \|B\|$.

8 If $A\boldsymbol{x} = \lambda\boldsymbol{x}$ then $\|A\boldsymbol{x}\|/\|\boldsymbol{x}\| = |\lambda|$ for that particular vector $\boldsymbol{x}$. When we maximize the ratio $\|A\boldsymbol{x}\|/\|\boldsymbol{x}\|$ over all vectors we get $\|A\| \ge |\lambda|$.

9 $A+B = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ has $\rho(A) = 0$ and $\rho(B) = 0$ but $\rho(A+B) = 1$.

The triangle inequality $\|A + B\| \le \|A\| + \|B\|$ fails for $\rho(A)$. $AB = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ has $\rho(AB) > \rho(A)\,\rho(B)$; thus $\rho(A) = \max|\lambda(A)|$ = spectral radius is not a norm.

10 (a) The condition number of A^{-1} is $\|A^{-1}\|\|(A^{-1})^{-1}\|$ which is $\|A^{-1}\|\|A\| = c(A)$.
(b) Since $A^{\mathrm{T}}A$ and AA^{T} have the same nonzero eigenvalues, A^{T} has the same norm as A.

11 Use the quadratic formula for $\lambda_{\max}/\lambda_{\min}$, which is $c = \sigma_{\max}/\sigma_{\min}$ since this $A = A^{\mathrm{T}}$ is positive definite:

$$c(A) = \left(1.00005 + \sqrt{(1.00005)^2 - .0001}\right) / \left(1.00005 - \sqrt{}\right) \approx 40,000.$$

12 $\det(2A)$ is not $2\det A$; $\det(A + B)$ is not always less than $\det A + \det B$; taking $|\det A|$ does not help. The only reasonable property is $\det AB = (\det A)(\det B)$. The condition number should not change when A is multiplied by 10.

13 The residual $\boldsymbol{b} - A\boldsymbol{y} = (10^{-7}, 0)$ is much smaller than $\boldsymbol{b} - A\boldsymbol{z} = (.0013, .0016)$. But $\boldsymbol{z}$ is much closer to the solution than $\boldsymbol{y}$.

14 $\det A = 10^{-6}$ so $A^{-1} = 10^3 \begin{bmatrix} 659 & -563 \\ -913 & 780 \end{bmatrix}$: $\|A\| > 1$, $\|A^{-1}\| > 10^6$, then $c > 10^6$.

15 $\boldsymbol{x} = (1,1,1,1,1)$ has $\|\boldsymbol{x}\| = \sqrt{5}, \|\boldsymbol{x}\|_1 = 5, \|\boldsymbol{x}\|_\infty = 1$. $\boldsymbol{x} = (.1,.7,.3,.4,.5)$ has $\|\boldsymbol{x}\| = 1, \|\boldsymbol{x}\|_1 = 2$ (sum), $\|\boldsymbol{x}\|_\infty = .7$ (largest).

16 $x_1^2+\cdots+x_n^2$ is not smaller than $\max(x_i^2)$ and not larger than $(|x_1|+\cdots+|x_n|)^2 = \|\boldsymbol{x}\|_1^2$. $x_1^2+\cdots+x_n^2 \leq n\,\max(x_i^2)$ so $\|\boldsymbol{x}\| \leq \sqrt{n}\|\boldsymbol{x}\|_\infty$. Choose $y_i = \text{sign}\,x_i = \pm 1$ to get $\|\boldsymbol{x}\|_1 = \boldsymbol{x}\cdot\boldsymbol{y} \leq \|\boldsymbol{x}\|\|\boldsymbol{y}\| = \sqrt{n}\|\boldsymbol{x}\|$. The vector $\boldsymbol{x} = (1,\ldots,1)$ has $\|\boldsymbol{x}\|_1 = \sqrt{n}\,\|\boldsymbol{x}\|$.

17 For the ℓ^∞ norm, the largest component of $\boldsymbol{x}$ plus the largest component of $\boldsymbol{y}$ is not less than $\|\boldsymbol{x}+\boldsymbol{y}\|_\infty$ = largest component of $\boldsymbol{x}+\boldsymbol{y}$.

For the ℓ^1 norm, each component has $|x_i+y_i| \leq |x_i|+|y_i|$. Sum on $i = 1$ to n: $\|\boldsymbol{x}+\boldsymbol{y}\|_1 \leq \|\boldsymbol{x}\|_1 + \|\boldsymbol{y}\|_1$.

18 $|x_1| + 2|x_2|$ is a norm but $\min(|x_1|,|x_2|)$ is not a norm. $\|\boldsymbol{x}\| + \|\boldsymbol{x}\|_\infty$ is a norm; $\|\boldsymbol{A}\boldsymbol{x}\|$ is a norm provided A is invertible (otherwise a nonzero vector has norm zero; for rectangular A we require independent columns to avoid $\|A\boldsymbol{x}\| = 0$).

19 $\boldsymbol{x}^{\mathrm{T}}\boldsymbol{y} = x_1y_1 + x_2y_2 + \cdots \leq (\max|y_i|)(|x_1|+|x_2|+\cdots) = ||\boldsymbol{x}||_1\,||\boldsymbol{y}||_\infty$.

20 With $\lambda_j = 2-2\cos(j\pi/n{+}1)$, the largest eigenvalue is $\lambda_n \approx 2+2 = 4$. The smallest is $\lambda_1 = 2-2\cos(\pi/n{+}1) \approx \left(\frac{\pi}{n+1}\right)^2$, using $2\cos\theta \approx 2-\theta^2$. So the condition number is $c = \lambda_{\max}/\lambda_{\min} \approx (4/\pi^2)\,n^2$, growing with n.

Problem Set 11.3, page 531

1 The iteration $\boldsymbol{x}_{k+1} = (I-A)\boldsymbol{x}_k + \boldsymbol{b}$ has $S = I$ and $T = I-A$ and $S^{-1}T = I-A$.

2 If $A\boldsymbol{x} = \lambda\boldsymbol{x}$ then $(I-A)\boldsymbol{x} = (1-\lambda)\boldsymbol{x}$. Real eigenvalues of $B = I-A$ have $|1-\lambda| < 1$ provided λ is between 0 and 2.

3 This matrix A has $I - A = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}$ which has $|\lambda| = 2$. The iteration diverges.

4 Always $\|AB\| \leq \|A\|\|B\|$. Choose $A = B$ to find $\|B^2\| \leq \|B\|^2$. Then choose $A = B^2$ to find $\|B^3\| \leq \|B^2\|\|B\| \leq \|B\|^3$. Continue (or use induction) to find $\|B^k\| \leq \|B\|^k$. Since $\|B\| \geq \max|\lambda(B)|$ it is no surprise that $\|B\| < 1$ gives convergence.

5 $A\boldsymbol{x}=\mathbf{0}$ gives $(S-T)\boldsymbol{x}=\mathbf{0}$. Then $S\boldsymbol{x}=T\boldsymbol{x}$ and $S^{-1}T\boldsymbol{x}=\boldsymbol{x}$. Then $\lambda=1$ means that the errors do not approach zero. We can't expect convergence when A is singular and $A\boldsymbol{x}=\boldsymbol{b}$ is unsolvable!

6 Jacobi has $S^{-1}T=\frac{1}{3}\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}$ with $|\lambda|_{\max}=\frac{1}{3}$. Small problem, fast convergence.

7 Gauss-Seidel has $S^{-1}T=\begin{bmatrix}0 & \frac{1}{3}\\ 0 & \frac{1}{9}\end{bmatrix}$ with $|\lambda|_{\max}=\frac{1}{9}$ which is $(|\lambda|_{\max}$ for Jacobi$)^{\mathbf{2}}$.

8 Jacobi has $S^{-1}T=\begin{bmatrix}a & \\ & d\end{bmatrix}^{-1}\begin{bmatrix}0 & -b\\ -c & 0\end{bmatrix}=\begin{bmatrix}0 & -b/a\\ -c/d & 0\end{bmatrix}$ with $|\lambda|=|bc/ad|^{1/2}$.

Gauss-Seidel has $S^{-1}T=\begin{bmatrix}a & 0\\ c & d\end{bmatrix}^{-1}\begin{bmatrix}0 & -b\\ 0 & 0\end{bmatrix}=\begin{bmatrix}0 & -b/a\\ 0 & -bc/ad\end{bmatrix}$ with $|\lambda|=|bc/ad|$.

So Gauss-Seidel is twice as fast to converge if $|\lambda|<1$ (or to explode if $|bc|>|ad|$).

9 Gauss-Seidel will converge for the $-1,2,-1$ matrix. $|\lambda|_{\max}=\cos^2\left(\frac{\pi}{n+1}\right)$ is given on page 527, together with the improvement from successive overrelaxation.

10 If the iteration gives all $x_i^{\text{new}}=x_i^{\text{old}}$ then the quantity in parentheses is zero, which means $A\boldsymbol{x}=\boldsymbol{b}$. For Jacobi change $\boldsymbol{x}^{\text{new}}$ on the right side to x^{old}.

11 $\boldsymbol{u}_k/\lambda_1^k=c_1\boldsymbol{x}_1+c_2\boldsymbol{x}_2(\lambda_2/\lambda_1)^k+\cdots+c_n\boldsymbol{x}_n(\lambda_n/\lambda_1)^k\to c_1\boldsymbol{x}_1$ if all ratios $|\lambda_i/\lambda_1|<1$. The largest ratio controls the rate of convergence (when k is large). $A=\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}$ has $|\lambda_2|=|\lambda_1|$ and no convergence.

12 The eigenvectors of A and also A^{-1} are $\boldsymbol{x}_1=(.75,.25)$ and $\boldsymbol{x}_2=(1,-1)$. The inverse power method converges to a multiple of $\boldsymbol{x}_2$, since $|1/\lambda_2|>|1/\lambda_1|$.

13 In the jth component of $A\boldsymbol{x}_1$, $\lambda_1\sin\frac{j\pi}{n+1}=2\sin\frac{j\pi}{n+1}-\sin\frac{(j-1)\pi}{n+1}-\sin\frac{(j+1)\pi}{n+1}$. The last two terms combine into $-2\sin\frac{j\pi}{n+1}\cos\frac{\pi}{n+1}$. Then $\lambda_1=2-2\cos\frac{\pi}{n+1}$.

14 $A = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$ produces $\boldsymbol{u}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\boldsymbol{u}_1 = \begin{bmatrix} 2 \\ -1 \end{bmatrix}$, $\boldsymbol{u}_2 = \begin{bmatrix} 5 \\ -4 \end{bmatrix}$, $\boldsymbol{u}_3 = \begin{bmatrix} 14 \\ -13 \end{bmatrix}$.

This is converging to the eigenvector direction $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$ with largest eigenvalue $\lambda = 3$.

Divide $\boldsymbol{u}_k$ by $\|\boldsymbol{u}_k\|$ to keep unit vectors.

15 $A^{-1} = \frac{1}{3}\begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ gives $\boldsymbol{u}_1 = \frac{1}{3}\begin{bmatrix} 2 \\ 1 \end{bmatrix}$, $\boldsymbol{u}_2 = \frac{1}{9}\begin{bmatrix} 5 \\ 4 \end{bmatrix}$, $\boldsymbol{u}_3 = \frac{1}{27}\begin{bmatrix} 14 \\ 13 \end{bmatrix} \to \boldsymbol{u}_\infty = \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix}$.

16 $R = Q^{\mathrm{T}}A = \begin{bmatrix} 1 & \cos\theta\sin\theta \\ 0 & -\sin^2\theta \end{bmatrix}$ and $A_1 = RQ = \begin{bmatrix} \cos\theta(1+\sin^2\theta) & -\sin^3\theta \\ \boldsymbol{-\sin^3\theta} & -\cos\theta\sin^2\theta \end{bmatrix}$.

17 If A is orthogonal then $Q = A$ and $R = I$. Therefore $A_1 = RQ = A$ again, and the "QR method" doesn't move from A. But shift A slightly and the method goes quickly to Λ.

18 If $A - cI = QR$ then $A_1 = RQ + cI = Q^{-1}(QR + cI)Q = Q^{-1}AQ$. No change in eigenvalues from the shift and shift back, because A_1 is similar to A.

19 Multiply $A\boldsymbol{q}_j = b_{j-1}\boldsymbol{q}_{j-1} + a_j\boldsymbol{q}_j + b_j\boldsymbol{q}_{j+1}$ by $\boldsymbol{q}_j^{\mathrm{T}}$ to find $\boldsymbol{q}_j^{\mathrm{T}}A\boldsymbol{q}_j = a_j$ (because the $\boldsymbol{q}$'s are orthonormal). The matrix form (multiplying by columns) is $AQ = QT$ where T is ***tridiagonal***. The entries down the diagonals of T are the a's and b's.

20 Theoretically the $\boldsymbol{q}$'s are orthonormal. In reality this important algorithm is not very stable. We must stop every few steps to reorthogonalize—or find another more stable way to orthogonalize the sequence $\boldsymbol{q}, A\boldsymbol{q}, A^2\boldsymbol{q}, \ldots$

21 If A is symmetric then $A_1 = Q^{-1}AQ = Q^{\mathrm{T}}AQ$ is also symmetric. $A_1 = RQ = R(QR)R^{-1} = RAR^{-1}$ has R and R^{-1} upper triangular, so A_1 cannot have nonzeros on a lower diagonal than A. If A is tridiagonal and symmetric then (by using symmetry for the upper part of A_1) the matrix $A_1 = RAR^{-1}$ is also tridiagonal.

22 From the last line of code, $\boldsymbol{q}_2$ is in the direction of $\boldsymbol{v} = A\boldsymbol{q}_1 - h_{11}\boldsymbol{q}_1 = A\boldsymbol{q}_1 - (\boldsymbol{q}_1^{\mathrm{T}}A\boldsymbol{q}_1)\boldsymbol{q}_1$. The dot product with $\boldsymbol{q}_1$ is zero. This is Gram-Schmidt with $A\boldsymbol{q}_1$ as the second input vector; we subtract from $A\boldsymbol{q}_1$ its projection onto the first vector $\boldsymbol{q}_1$.

Note The three lines after the short "pseudocodes" describe two key properties of conjugate gradients—the residuals $\boldsymbol{r}_k = \boldsymbol{b} - A\boldsymbol{x}_k$ are orthogonal and the search directions are A-orthogonal ($\boldsymbol{d}_i^{\mathrm{T}} A \boldsymbol{d}_k = 0$). Then each new approximation $\boldsymbol{x}_{k+1}$ is the **closest vector to $\boldsymbol{x}$** among all combinations of $\boldsymbol{b}, A\boldsymbol{b} \ldots, A^k\boldsymbol{b}$. Ordinary iteration $S\boldsymbol{x}_{k+1} = T\boldsymbol{x}_k + \boldsymbol{b}$ does *not* find this best possible combination $\boldsymbol{x}_{k+1}$.

23 The solution is straightforward and important. Since $H = Q^{-1}AQ = Q^{\mathrm{T}}AQ$ is ***symmetric*** if $A = A^{\mathrm{T}}$, and since H has only one lower diagonal by construction, then H has only ***one upper diagonal***: H is tridiagonal and all the recursions in Arnoldi's method have only 3 terms.

24 $H = Q^{-1}AQ$ is ***similar*** to A, so H has the same eigenvalues as A (at the end of Arnoldi). When Arnoldi is stopped sooner because the matrix size is large, the eigenvalues of H_k (called ***Ritz values***) are close to eigenvalues of A. This is an important way to compute approximations to λ for large matrices.

25 In principle the conjugate gradient method converges in 100 (or 99) steps to the exact solution $\boldsymbol{x}$. But it is slower than elimination and its all-important property is to give good approximations to $\boldsymbol{x}$ much sooner. (Stopping elimination part way leaves you nothing.) The problem asks how close $\boldsymbol{x}_{10}$ and $\boldsymbol{x}_{20}$ are to $\boldsymbol{x}_{100}$, which equals $\boldsymbol{x}$ except for roundoff errors.

26 $A = \begin{bmatrix} 1 & 1 \\ 0 & 1.1 \end{bmatrix}$ has $A^n = \begin{bmatrix} 1 & q \\ 0 & (1.1)^n \end{bmatrix}$ with $q = 1 + 1.1 + \cdots + (1.1)^{n-1} = (1.1^n - 1)/(1.1 - 1) \approx \mathbf{10\,(1.1)^{\boldsymbol{n}}}$. So the growing part of A^n is $(1.1)^n \begin{bmatrix} 0 & 10 \\ 0 & 1 \end{bmatrix}$ with $||A^n|| \approx \sqrt{101}$ times 1.1^n for larger n.