STAT 309: MATHEMATICAL COMPUTATIONS I
FALL 2013
LECTURE 11

1. COMPUTING THE QR FACTORIZATION

- there are two common ways to compute the QR decomposition:
  - using *Householder matrices*, developed by Alston S. Householder
  - using *Givens rotations*, also known as *Jacobi rotations*, used by Wallace Givens and originally invented by Jacobi for use with in solving the symmetric eigenvalue problem in 1846
  - the Gram-Schmidt or modfied Gram-Schmidt orthogonalization discussed in previous lecture works in principle but has numerical stability issues and are not usually used
- roughly speaking, Gram-Schmidt applies a sequence of triangular matrices to orthogonalize A (i.e., transform A into an orthogonal matrix Q),
$$AR_1^{-1}R_2^{-1}\cdots R_{n-1}^{-1} = Q$$
whereas Householder and Givens QR apply a sequence of orthogonal matrices to triangularize A (i.e., transform A into an upper triangular matrix R),
$$Q_{n-1}^{\mathsf{T}}\cdots Q_2^{\mathsf{T}}Q_1^{\mathsf{T}}A = R$$
- orthogonal transformations are highly desirable in algorithms as they preserve lengths and therefore do not blow up the errors present at every stage of the computation

2. ORTHOGONALIZATION USING GIVENS ROTATIONS

- we illustrate the process in the case where A is a 2×2 matrix
- in Gaussian elimination, we compute $L^{-1}A = U$ where L^{-1} is unit lower triangular and U is upper triangular, specifically,
$$\begin{bmatrix} 1 & 0 \\ m_{21} & 1 \end{bmatrix}\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \begin{bmatrix} a_{11}^{(2)} & a_{12}^{(2)} \\ 0 & a_{22}^{(2)} \end{bmatrix}, \quad m_{21} = -\frac{a_{21}}{a_{11}}$$
- by contrast, the QR decomposition takes the form
$$\begin{bmatrix} \gamma & \sigma \\ -\sigma & \gamma \end{bmatrix}\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} \\ 0 & r_{22} \end{bmatrix}$$
where $\gamma^2 + \sigma^2 = 1$
- from the relationship $-\sigma a_{11} + \gamma a_{21} = 0$ we obtain
$$\begin{aligned} \gamma a_{21} &= \sigma a_{11} \\ \gamma^2 a_{21}^2 &= \sigma^2 a_{11}^2 = (1-\gamma^2)a_{11}^2 \end{aligned}$$
which yields
$$\gamma = \pm\frac{a_{11}}{\sqrt{a_{21}^2 + a_{11}^2}}$$
- it is conventional to choose the + sign

Date: November 8, 2013, version 1.1. Comments, bug reports: `lekheng@galton.uchicago.edu`.

- then, we obtain

$$\sigma^2 = 1 - \gamma^2 = 1 - \frac{a_{11}^2}{a_{21}^2 + a_{11}^2} = \frac{a_{21}^2}{a_{21}^2 + a_{11}^2},$$

  or

$$\sigma = \pm \frac{a_{21}}{\sqrt{a_{21}^2 + a_{11}^2}}$$

- again, we choose the + sign
- as a result, we have

$$r_{11} = a_{11} \frac{a_{11}}{\sqrt{a_{21}^2 + a_{11}^2}} + a_{21} \frac{a_{21}}{\sqrt{a_{21}^2 + a_{11}^2}} = \sqrt{a_{21}^2 + a_{11}^2}$$

- the matrix

$$Q^\mathsf{T} = \begin{bmatrix} \gamma & \sigma \\ -\sigma & \gamma \end{bmatrix}$$

  is called a *Givens rotation*
- it is called a rotation because it is orthogonal, and therefore length-preserving, and also because there is an angle θ such that $\sin\theta = \sigma$ and $\cos\theta = \gamma$, and its effect is to rotate a vector through the angle θ
- in particular,

$$\begin{bmatrix} \gamma & \sigma \\ -\sigma & \gamma \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \rho \\ 0 \end{bmatrix}$$

  where $\rho = \sqrt{\alpha^2 + \beta^2}$, $\alpha = \rho\cos\theta$ and $\beta = \rho\sin\theta$
- it is easy to verify that the product of two rotations is itself a rotation
- now, in the case where A is an $n \times n$ matrix, suppose that we are given the vector

$$\begin{bmatrix} \times & \cdots & \times & \alpha & \times & \cdots & \times & \beta & \times & \cdots & \times \end{bmatrix}^\mathsf{T} \in \mathbb{R}^n,$$

  then

$$\begin{bmatrix} 1 & & & & & & & & & & \\ & \ddots & & & & & & & & & \\ & & 1 & & & & & & & & \\ & & & \gamma & & & & \sigma & & & \\ & & & & 1 & & & & & & \\ & & & & & \ddots & & & & & \\ & & & & & & 1 & & & & \\ & & & -\sigma & & & & \gamma & & & \\ & & & & & & & & 1 & & \\ & & & & & & & & & \ddots & \\ & & & & & & & & & & 1 \end{bmatrix} \begin{bmatrix} \times \\ \vdots \\ \times \\ \alpha \\ \times \\ \vdots \\ \times \\ \beta \\ \times \\ \vdots \\ \times \end{bmatrix} = \begin{bmatrix} \times \\ \vdots \\ \times \\ \rho \\ \times \\ \vdots \\ \times \\ 0 \\ \times \\ \vdots \\ \times \end{bmatrix}$$

- so, in order to transform A into an upper triangular matrix R, we can find a product of rotations Q such that $Q^\mathsf{T} A = R$
- it is easy to see that $O(n^2)$ rotations are required

3. Orthogonalization using Householder reflections

- it is natural to ask whether we can introduce more zeros with each orthogonal rotation and to that end, we examine *Householder reflections*
- consider a matrix of the form $P = I - \tau\mathbf{u}\mathbf{u}^\mathsf{T}$, where $\mathbf{u} \neq \mathbf{0}$ and τ is a nonzero constant
- a P that has this form is called a *symmetric rank*-1 *change* of I
- can we choose τ so that P is also orthogonal?

- from the desired relation $P^\top P = I$ we obtain

$$\begin{aligned} P^\top P &= (I - \tau \mathbf{u}\mathbf{u}^\top)^\top (I - \tau \mathbf{u}\mathbf{u}^\top) \\ &= I - 2\tau \mathbf{u}\mathbf{u}^\top + \tau^2 \mathbf{u}\mathbf{u}^\top \mathbf{u}\mathbf{u}^\top \\ &= I - 2\tau \mathbf{u}\mathbf{u}^\top + \tau^2 (\mathbf{u}^\top \mathbf{u}) \mathbf{u}\mathbf{u}^\top \\ &= I - (\tau^2 \mathbf{u}^\top \mathbf{u} - 2\tau) \mathbf{u}\mathbf{u}^\top \\ &= I + \tau(\tau \mathbf{u}^\top \mathbf{u} - 2) \mathbf{u}\mathbf{u}^\top \end{aligned}$$

- it follows that if $\tau = 2/\mathbf{u}^\top \mathbf{u}$, then $P^\top P = I$ for any nonzero $\mathbf{u}$
- without loss of generality, we can stipulate that $\mathbf{u}^\top \mathbf{u} = 1$, and therefore P takes the form $P = I - 2\mathbf{v}\mathbf{v}^\top$, where $\mathbf{v}^\top \mathbf{v} = 1$
- why is the matrix P called a reflection?
- this is because for any nonzero vector $\mathbf{x}$, $P\mathbf{x}$ is the reflection of $\mathbf{x}$ across the hyperplane that is normal to $\mathbf{v}$
- for example, consider the 2×2 case and set $\mathbf{v} = \begin{bmatrix} 1 & 0 \end{bmatrix}^\top$ and $\mathbf{x} = \begin{bmatrix} 1 & 2 \end{bmatrix}^\top$, then

$$P = I - 2\mathbf{v}\mathbf{v}^\top = I - 2 \begin{bmatrix} 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - 2 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$$

  and therefore

$$P\mathbf{x} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$$

- now, let $\mathbf{x}$ be any vector, we wish to construct P so that $P\mathbf{x} = \alpha \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}^\top = \alpha \mathbf{e}_1$ for some α
- from the relations

$$\|P\mathbf{x}\|_2 = \|\mathbf{x}\|_2, \qquad \|\alpha \mathbf{e}_1\|_2 = |\alpha| \|\mathbf{e}_1\|_2 = |\alpha|$$

  we obtain $\alpha = \pm \|\mathbf{x}\|_2$
- to determine P, we observe that

$$\mathbf{x} = P^{-1}(\alpha \mathbf{e}_1) = \alpha P \mathbf{e}_1 = \alpha (I - 2\mathbf{v}\mathbf{v}^\top)\mathbf{e}_1 = \alpha(\mathbf{e}_1 - 2\mathbf{v}\mathbf{v}^\top \mathbf{e}_1) = \alpha(\mathbf{e}_1 - 2\mathbf{v} v_1)$$

  which yields the system of equations

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \alpha \begin{bmatrix} 1 - 2v_1^2 \\ -2v_1 v_2 \\ \vdots \\ -2v_1 v_n \end{bmatrix}$$

- from the first equation $x_1 = \alpha(1 - 2v_1^2)$ we obtain

$$v_1 = \pm \sqrt{\frac{1}{2}\left(1 - \frac{x_1}{\alpha}\right)}$$

- for $i = 2, \ldots, n$, we have

$$v_i = -\frac{x_i}{2\alpha v_1}$$

- it is best to choose α to have the opposite sign of x_1 to avoid cancellation in v_1
- it is conventional to choose the $+$ sign for α if $x_1 = 0$
- note that the matrix P is never formed explicitly: for any vector $\mathbf{b}$, the product $P\mathbf{b}$ can be computed as follows

$$P\mathbf{b} = (I - 2\mathbf{v}\mathbf{v}^\top)\mathbf{b} = \mathbf{b} - 2(\mathbf{v}^\top \mathbf{b})\mathbf{v}$$

- this process requires only $O(2n)$ operations

- it is easy to see that we can represent P simply by storing only $\mathbf{v}$
- we showed how a Householder reflection of the form $P = I - 2\mathbf{u}\mathbf{u}^\top$ could be constructed so that given a vector $\mathbf{x}$, $P\mathbf{x} = \alpha\mathbf{e}_1$
- now, suppose that that $\mathbf{x} = \mathbf{a}_1$ is the first column of a matrix A, then we construct a Householder reflection $H_1 = I - 2\mathbf{u}_1\mathbf{u}_1^\top$ such that $H\mathbf{x} = \alpha\mathbf{e}_1$, and we have

$$A^{(2)} = H_1 A = \begin{bmatrix} r_{11} & & & \\ 0 & & & \\ \vdots & \mathbf{a}_2^{(2)} & \cdots & \mathbf{a}_n^{(2)} \\ 0 & & & \end{bmatrix}$$

  where we denote the constant α by r_{11}, as it is the $(1,1)$ element of the updated matrix $A^{(2)}$
- now, we can construct H_2 such that

$$H_2\mathbf{a}^{(2)} = \begin{bmatrix} a_{12}^{(2)} \\ r_{22} \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad u_{12} = 0, \quad H_2 = \begin{bmatrix} 1 & 0 \\ 0 & \\ \vdots & h_{ij} \\ 0 & \end{bmatrix}$$

- note that the first column of $A^{(2)}$ is unchanged by H_2
- continuing this process, we obtain

$$H_{n-1} \cdots H_1 A = A^{(n)} = R$$

  where R is an upper triangular matrix
- we have thus factored $A = QR$, where $Q = H_1 H_2 \cdots H_{n-1}$ is an orthogonal matrix
- note that

$$A^\top A = R^\top Q^\top Q R = R^\top R,$$

  and thus R is the Cholesky factor of $A^\top A$ (we will discus Cholesky factorization next week)

4. GIVENS ROTATIONS VERSUS HOUSEHOLDER REFLECTIONS

- we showed how to construct Givens rotations in order to rotate two elements of a column vector so that one element would be zero, and that approximately $n^2/2$ such rotations could be used to transform A into an upper triangular matrix R
- because each rotation only modifies two rows of A, it is possible to interchange the order of rotations that affect different rows, and thus apply sets of rotations in parallel
- this is the main reason why Givens rotations can be preferable to Householder reflections
- other reasons are that they are easy to use when the QR factorization needs to be updated as a result of adding a row to A or deleting a column of A
- Givens rotations are also more efficient when A is sparse

5. COMPUTING THE COMPLETE ORTHOGONAL FACTORIZATION

- we first seek a decomposition of the form $A = QR\Pi$ where the permutation matrix Π is chosen so that the diagonal elements of R are maximized at each stage
- specifically, suppose

$$H_1 A = \begin{bmatrix} r_{11} & \times & \cdots & \times \\ 0 & \times & \cdots & \times \\ \vdots & \vdots & & \vdots \\ 0 & \times & \cdots & \times \end{bmatrix}, \quad r_{11} = \|\mathbf{a}_1\|_2$$

- so, we choose Π_1 so that $\|\mathbf{a}_1\|_2 \geq \|\mathbf{a}_j\|_2$ for $j \geq 2$
- for Π_2, look at the lengths of the columns of the submatrix; we don't need to recompute the lengths each time, because we can update by subtracting the square of the first component from the square of the total length
- eventually, we get

$$Q \begin{bmatrix} R & S \\ 0 & 0 \end{bmatrix} \Pi_1 \cdots \Pi_r = A$$

  where R is upper triangular
- suppose

$$A = Q \begin{bmatrix} R & S \\ 0 & 0 \end{bmatrix} \Pi$$

  where R is upper triangular, then

$$A^\mathsf{T} = \Pi^\mathsf{T} \begin{bmatrix} R^\mathsf{T} & 0 \\ S^\mathsf{T} & 0 \end{bmatrix} Q^\mathsf{T}$$

  where R^T is lower triangular
- we apply Householder reflections so that

$$H_i \cdots H_2 H_1 \begin{bmatrix} R^\mathsf{T} & 0 \\ S^\mathsf{T} & 0 \end{bmatrix} = \begin{bmatrix} U & 0 \\ 0 & 0 \end{bmatrix}$$

- then

$$A^\mathsf{T} = Z^\mathsf{T} \begin{bmatrix} U & 0 \\ 0 & 0 \end{bmatrix} Q^\mathsf{T}$$

  where $Z = H_i \cdots H_1 \Pi$