Lecture 11: Gram-Schmidt Orthogonalization

June 18, 2025

Outline

QR factorization via Gram-Schmidt

- **2** Upper Triangular Matrix R
- 2 Modified Gram-Schmidt
- Omplexity of Gram-Schmidt

Gram-Schmidt Orthogonalization - Introduction

- We have already seen a variation of Gram-Schmidt orthogonalization in constructing A-orthogonal search directions for the conjugate gradient method (see Lecture 09).
- The same general idea applies here to construct the QR factorization:
 - Use columns of A as proposed vectors to be orthogonalized into Q.
 - Build each new vector q_j by orthogonalizing a_j with respect to all previous q vectors, { q₁, q₂,..., q_{j-1} }, and then normalizing it.

 For a vector a_j, we can orthogonalize it against all previous vectors q_i for i = 1,..., j − 1, using

$$v_j = a_j - (q_1^T a_j)q_1 - (q_2^T a_j)q_2 - \dots - (q_{j-1}^T a_j)q_{j-1}.$$
 (1)

- This removes *a_j*'s components in the orthogonal directions constructed so far.
- We can then directly normalize, giving us

$$q_j = \frac{v_j}{\|v_j\|_2}.$$

• Equation (1) was derived previously in Lecture 09 assuming the form $v_j = a_j + \sum_{i=1}^{j-1} \beta_i q_i$.

- To recap, the coefficients β_j were then derived as follows.
- Since we want v_j to be orthogonal to all previous q_i's, we get the equation (for some 1 ≤ k ≤ j − 1):

$$0 = q_k^T v_j,$$

= $q_k^T a_j + \sum_{i=1}^{j-1} \beta_i (q_k^T q_i),$
= $q_k^T a_j + \beta_k (q_k^T q_k).$

Since q_k^Tq_k = 1 we have that β_k = -q_k^Ta_j giving the equation in (1)

$$v_j = a_j - \sum_{i=1}^{j-1} (q_i^T a_j) q_i$$

- Consider the following example in 2D.
- We are given a_2 and the (previous) unit vector q_1 .
- We want to find q_2 that is orthonormal to q_1 .



- We apply the following steps:
 - Orthogonalize using $v_2 = a_2 (q_1^T a_2)q_1$,
 - 2 Normalize to get $q_2 = \frac{v_2}{\|v_2\|_2}$.
- So we have that

$$q_2 = \frac{a_2 - (q_1^T a_2)q_1}{\|a_2 - (q_1^T a_2)q_1\|_2}$$

• But what are the entries of R in our QR factorization?

To get the entries of R consider the general form of the $\ensuremath{\mathsf{QR}}$ factorization

$$\begin{bmatrix} | & | & & | \\ a_1 & a_2 & \dots & a_n \\ | & | & & | \end{bmatrix} = \begin{bmatrix} | & | & & | \\ q_1 & q_2 & \dots & q_n \\ | & | & & | \end{bmatrix} \begin{bmatrix} r_{11} & \dots & r_{1n} \\ & \ddots & \vdots \\ & & r_{nn} \end{bmatrix}.$$

Written out componentwise we have

$$\begin{array}{rcl} a_{1} = r_{11}q_{1} & \Rightarrow & q_{1} = \frac{a_{1}}{r_{11}}, \\ a_{2} = r_{12}q_{1} + r_{22}q_{2} & \Rightarrow & q_{2} = \frac{a_{2}-r_{12}q_{1}}{r_{22}}, \\ a_{3} = r_{13}q_{1} + r_{23}q_{2} + r_{33}q_{3} & \Rightarrow & q_{3} = \frac{a_{3}-r_{13}q_{1}-r_{23}q_{2}}{r_{33}}, \\ \vdots & & \vdots \\ a_{n} = r_{1n}q_{1} + r_{2n}q_{2} + \ldots + r_{nn}q_{n} & \Rightarrow & q_{n} = \frac{a_{n}-\sum_{i=1}^{n-1}r_{in}q_{i}}{r_{nn}}. \end{array}$$

- Now we compare the *q_i* above with the result from our Gram-Schmidt orthogonalization expressions.
- For the 2D example above we have

Gram-Schmidt:
$$q_2 = \frac{a_2 - (q_1^T a_2)q_1}{\|a_2 - (q_1^T a_2)q_1\|_2}$$
,
Factorization: $q_2 = \frac{a_2 - r_{12}q_1}{r_{22}}$.

So

$$egin{array}{rll} r_{12} &=& q_1^T a_2, ext{ and } \ r_{22} &=& \|a_2 - (q_1^T a_2) q_1\|_2. \end{array}$$

• In general (higher dimensions), the entries of *R* can be written as

$$r_{ij} = (q_i^T a_j),$$

$$r_{jj} = \left\| a_j - \sum_{i=1}^{j-1} r_{ij} q_i \right\|_2$$

The off-diagonal entries r_{ij} correspond to the lengths of components of a_j in previous directions q₁,..., q_{j-1}.
 The diagonal entries r_{jj} correspond to the length of v_j, required to normalize to q_i.

- The classic Gram-Schmidt (CGS) algorithm for QR factorization is given in Algorithm 12.
- Note that the classic Gram-Schmidt is numerically unstable.
- That is, it is sensitive to round off error (and can even yield non-orthogonal q's).

Algorithm 1 : Gram-Schmidt Algorithm (Classic)

for
$$j = 1, 2, ..., n$$

 $v_j = a_j$ \triangleright get next column
for $i = 1, 2, ..., j - 1$ \triangleright Orthogonalize
 $r_{ij} = q_i^T a_j$
 $v_j = v_j - r_{ij}q_i$
end for
 $r_{jj} = ||v_j||_2$ \triangleright Normalize
 $q_j = v_j/r_{jj}$
end for

We can alter the classic Gram-Schmidt algorithm in the inner loop for better numerical stability.

Algorithm 2 : QR Factorization With Modified Gram-Schmidt Algorithm

1: **for** i = 1 : n▷ Get next column 2: $v_i = a_i$ 3: end for 4: for j = 1 : n $r_{ii} = \|v_i\|_2$ 5: $q_j = \left(\frac{1}{r_{ij}}\right) v_j$ 6: ▷ Normalize **for** k = j + 1 : n7: \triangleright Orthogonalize $r_{ik} = q_i^T v_k$ 8. 9. $v_k = v_k - r_{ik}q_i$ end for $10 \cdot$ 11: end for

The modified Gram-Schmidt algorithm (see Algorithm 13) gives an identical result as the classic Gram-Schmidt algorithm **in exact arithmetic**.

In floating point arithmetic, the modified Algorithm 13 is more numerically stable than the classic Algorithm 12. Think carefully about these pseudocodes.

- In classical Gram-Schmidt, we take each vector, one at a time, and make it orthogonal to all previous vectors.
- In modified Gram-Schmidt, we take each vector, and modify all forthcoming vectors to be orthogonal to it.

Once you argue this way, it is clear that both methods are performing the same operations, and are mathematically equivalent.

- But, importantly, modified Gram-Schmidt suffers from round-off instability to a significantly lesser degree.
- This can be explained, in part, from the formulas for Gram-Schmidt, without QR-factorization:

$$CGS: v_k = a_k - \left(q_j^T a_k\right) q_j, \text{ and}$$
$$MGS: v_k = v_k - \left(q_j^T v_k\right) q_j.$$

If an error is made in computing q_2 in CGS, so that $q_1^T q_2 = \delta$ is small, but non-zero, this will not be corrected for in any of the computations that follow:

$$v_{3} = a_{3} - (q_{1}^{T}a_{3})q_{1} - (q_{2}^{T}a_{3})q_{2},$$

$$q_{2}^{T}v_{3} = q_{2}^{T}a_{3} - q_{2}^{T}(q_{1}^{T}a_{3})q_{1} - q_{2}^{T}(q_{2}^{T}a_{3})q_{2},$$

$$= q_{2}^{T}a_{3} - (q_{1}^{T}a_{3})\underbrace{q_{2}^{T}q_{1}}_{=\delta} - (q_{2}^{T}a_{3})\underbrace{q_{2}^{T}q_{2}}_{=1},$$

$$= q_{2}^{T}a_{3} - (q_{1}^{T}a_{3})\delta - (q_{2}^{T}a_{3})$$

$$= -(q_{1}^{T}a_{3})\delta.$$

Similarly,

$$v_{3} = a_{3} - (q_{1}^{T}a_{3})q_{1} - (q_{2}^{T}a_{3})q_{2},$$

$$q_{1}^{T}v_{3} = q_{1}^{T}a_{3} - q_{1}^{T}(q_{1}^{T}a_{3})q_{1} - q_{1}^{T}(q_{2}^{T}a_{3})q_{2},$$

$$= q_{1}^{T}a_{3} - (q_{1}^{T}a_{3})\underbrace{q_{1}^{T}q_{1}}_{=1} - (q_{2}^{T}a_{3})\underbrace{q_{1}^{T}q_{2}}_{=\delta},$$

$$= q_{1}^{T}a_{3} - (q_{1}^{T}a_{3}) - (q_{2}^{T}a_{3})\delta$$

$$= -(q_{2}^{T}a_{3})\delta.$$

We see that v_3 is not orthogonal to either of q_1 or q_2 .

On the other hand, assume the same for MGS (n = 3): $q_1^T q_2 = \delta$ is small, but non-zero. Let's examine how the third vector v_3 changes:

nitially,
$$v_3^{(0)} = a_3$$
.
• $\underline{j} = 1, k = 3$

$$\begin{array}{rcl} r_{13} & = & q_1^T v_3^{(0)} \\ v_3^{(1)} & = & v_3^{(0)} - r_{13} q_1 \\ & = & v_3^{(0)} - (q_1^T v_3^{(0)}) q_1 \end{array}$$

$$\begin{array}{rcl} r_{23} & = & q_2^T v_3^{(1)} \\ v_3^{(2)} & = & v_3^{(1)} - r_{23} q_2 \\ & = & v_3^{(1)} - (q_2^T v_3^{(1)}) q_2 \end{array}$$

Hence we obtain

$$q_{2}^{T}v_{3}^{(2)} = q_{2}^{T}v_{3}^{(1)} - q_{2}^{T}(q_{2}^{T}v_{3}^{(1)})q_{2}$$

$$= q_{2}^{T}v_{3}^{(1)} - (q_{2}^{T}v_{3}^{(1)})\underbrace{q_{2}^{T}q_{2}}_{=1}$$

$$= 0$$

$$q_{1}^{T}v_{3}^{(2)} = q_{1}^{T}v_{3}^{(1)} - q_{1}^{T}(q_{2}^{T}v_{3}^{(1)})q_{2}$$

$$= q_{1}^{T}v_{3}^{(1)} - (q_{2}^{T}v_{3}^{(1)})\underbrace{q_{1}^{T}q_{2}}_{=\delta}$$

$$= q_{1}^{T}v_{3}^{(1)} - q_{2}^{T}v_{3}^{(1)}\delta$$

(2)

(3)

Computing each of the terms on line (3) gives

$$\begin{aligned} q_1^T v_3^{(1)} &= q_1^T [v_3^{(0)} - (q_1^T v_3^{(0)})q_1] \\ &= q_1^T v_3^{(0)} - q_1^T (q_1^T v_3^{(0)})q_1] \\ &= q_1^T v_3^{(0)} - (q_1^T v_3^{(0)}) \underbrace{q_1^T q_1}_{=1} \\ &= q_1^T v_3^{(0)} - (q_1^T v_3^{(0)}) \underbrace{q_1^T q_1}_{=1} \\ &= 0, \text{ and} \\ q_2^T v_3^{(1)} &= q_2^T [v_3^{(0)} - (q_1^T v_3^{(0)})q_1] \\ &= q_2^T v_3^{(0)} - q_2^T (q_1^T v_3^{(0)})q_1 \\ &= q_2^T v_3^{(0)} - (q_1^T v_3^{(0)}) \underbrace{q_2^T q_1}_{=\delta} \\ &= q_2^T v_3^{(0)} - (q_1^T v_3^{(0)}) \delta \end{aligned}$$

Putting it all together gives

$$q_{1}^{T}v_{3}^{(2)} = q_{1}^{T}v_{3}^{(1)} - q_{2}^{T}v_{3}^{(1)}\delta$$

= $0 - [q_{2}^{T}v_{3}^{(0)} - (q_{1}^{T}v_{3}^{(0)})\delta]\delta$
= $-q_{2}^{T}v_{3}^{(0)}\delta + (q_{1}^{T}v_{3}^{(0)})\delta^{2}$ (4)

- Recall that $v_3 = v_3^{(2)}$ is the final form of the third vector (before normalization).
- Let's check orthogonality, assuming no more errors are made.
- First for q_2 :

$$q_2^T v_3^{(2)} = 0$$
, from equation (2).

So, we perserve orthogonality to q_2 .

• Second for *q*₁:

$$q_1^T v_3^{(2)} = -q_2^T v_3^{(0)} \delta + (q_1^T v_3^{(0)}) \delta^2$$
, from equation (4).

Summarized Comparison of CGS versus MGS

Inner Product	CGS	MGS
$q_2^T v_3$	$-(q_1^T a_3)\delta$	0
$q_1^T v_3$	$-(q_2^T a_3)\delta$	$-q_2^T v_3^{(0)} \delta + (q_1^T v_3^{(0)}) \delta^2$
		$= -q_2^T a_3 \delta + q_1^T a_3 \delta^2$
		$v_3^{(0)} = a_3$

Remarks:

- $\begin{tabular}{ll} \hline \begin{tabular}{ll} \hline \end{tabular} \\ \hline \$
- **2** Because of the opposite signs of the terms involved, it is likely, but not guaranteed, that $|-q_2^T a_3 \delta + q_1^T a_3 \delta^2| \le |-(q_2^T a_3)\delta|$.
- **③** The error in $q_1^T v_3$ is likely no worse than in CGS, but we have eliminated the errors in $q_2^T v_3$, an improvement.

(Reference: https://www.math.uci.edu/~ttrogdon/105A/ html/Lecture23.html)

Explanation for the Computation of r_{jk}

- Note that the Modified Gram-Schmidt algorithm uses $r_{jk} = q_j^T v_k$, where one might instead expect $r_{jk} = q_j^T a_k$.
- Here I will prove that these two choices must always yield identical results.
- Use the notation of the algorithm throughout the explanation.

- Let $1 \le j \le n$ be arbitrary.
- Let $j + 1 \le k \le n$ be arbitrary.
- I claim that $q_j^T v_k = q_j^T a_k$.
- Note that v_k is updated once per (outer) *j*-loop.
- Since that update takes place AFTER the computation of the $q_j^T v_k$ inner product, at the time of that computation,

$$\begin{aligned} \mathbf{v}_k &= \mathbf{a}_k - \sum_{\ell=1}^{j-1} r_{\ell k} q_\ell, \text{ so that} \\ q_j^T \mathbf{v}_k &= q_j^T \left[\mathbf{a}_k - \sum_{\ell=1}^{j-1} r_{\ell k} q_\ell \right] \\ &= q_j^T \mathbf{a}_k - \sum_{\ell=1}^{j-1} r_{\ell k} \underbrace{q_j^T q_\ell}_{=0, \text{ since } \ell < j} \\ &= q_j^T \mathbf{a}_k, \text{ as claimed.} \end{aligned}$$

 The example below computes the QR factorization of a random matrix with hugely varying magnitudes of r_{ii} (diagonal) entries.



- The blue points are the result using the classic Gram-Schmidt algorithm.
- It runs out of accuracy at around $\sqrt{E_{\text{machine}}}$.
- The orange points are the result with the **modified** Gram-Schmidt algorithm.
- The modified algorithm runs out of accuracy at around $E_{
 m machine}$.
- For more information see Lecture 9, Experiment 2 in Trefethen & Bau.

Another exercise is to find the QR factorization of

$$A = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}$$

via Gram-Schmidt orthogonalization. The solution to this problem is

$$Q = \begin{bmatrix} \frac{\sqrt{2}}{2} & \frac{\sqrt{3}}{3} & -\frac{\sqrt{6}}{6} \\ 0 & \frac{\sqrt{3}}{3} & \frac{\sqrt{6}}{3} \\ \frac{\sqrt{2}}{2} & -\frac{\sqrt{3}}{3} & \frac{\sqrt{6}}{6} \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} \sqrt{2} & \sqrt{2} & \frac{\sqrt{2}}{2} \\ 0 & \sqrt{3} & 0 \\ 0 & 0 & \frac{\sqrt{6}}{2} \end{bmatrix}$$

See Lecture Notes for the details.

Gram-Schmidt Orthogonalization - Complexity of Gram-Schmidt

- The inner *i*-loop involves
 - $r_{ij} = q_i^T a_j$ (CGS) or $q_i^T v_j$ (MGS) $\Rightarrow m$ (scalar) multiplications and m - 1 additions for the inner products,
 - v_j = v_j − r_{ij}q_i (CGS) or v_j = v_j − r_{jk}q_j (MGS) ⇒ m multiplications, and m subtractions.
- Hence, the flops per inner loop $\approx 4m$.
- An approximation of the total flops is therefore

$$\sum_{j=1}^{n} \sum_{i=1}^{j-1} 4m = 4m \sum_{j=1}^{n} (j-1),$$

$$\approx 4m \sum_{j=1}^{n} = 4m \frac{n(n+1)}{2},$$

$$\approx 2mn^{2}.$$

Gram-Schmidt Orthogonalization - Complexity of Gram-Schmidt

- If the matrix is square (m = n), flops(Gram-Schmidt) = $2n^3 + O(n^2) \approx 3 \times \text{flops}(LU)$.
- We could use QR factorization to solve linear systems also, but at a cost $3 \times$ greater than LU factorization.
- As was pointed out earlier, if A is square and A = QR, then solving Ax = b for x is equivalent to solving Rx = Q^Tb for x.