***Class Notes*, Math 554, Autumn 2012**

Lecture XXI: The Gram-Schmidt Process and the QR factorization; Least Squares

The Gram-Schmidt algorithm started on the columns of a matrix A will yield the *QR factorization.*

Proposition 1. *Given $A \in \mathcal{M}_{m,n}(\mathbb{C})$, $m \geq n$, there exists a unitary matrix $Q \in \mathcal{M}_m(\mathbb{C})$ and an upper triangular matrix $R \in \mathcal{M}_{m,n}(\mathbb{C})$ such that $A = QR$. (Note that for a rectangular matrix to be upper triangular it means that all entries with indices (i, j), $i > j$ are 0.)*

Proof. Suppose first that the columns of A are independent. The first n columns of Q, $\tilde{Q}$, will be obtained as a result of the Gram-Schmidt process on A. The same process also yields an upper triangular matrix $\tilde{R}$ which is square $(n]timesn)$.

Denote $Q = [\tilde{Q}, Q']$ a completion of $\tilde{Q}$ to full orthonormal basis and by $R = \begin{bmatrix} \tilde{R} \\ 0 \end{bmatrix}$ the completion of R to $m \times n$ by adding $m - n$ rows of zeros.

Then $A = QR$.

Now, if the columns of A are dependent, one can continue the Gram-Schmidt process by setting appropriate qs to 0. For example, if a_k is a linear combination of the previous columns, which will be discovered as p_k will be 0, set $a_k = 0$ and also all $r_{kj} = 0-$ for $j \geq k$, and continue. The result will once again be a matrix $\tilde{Q}$ with orthogonal columns, some of which are unit-length and some which are 0, as well as an upper triangular matrix $\tilde{R}$ with some zero rows. Eliminate all zero columns and correspdondingly zero rows from $\tilde{Q}$ and $\tilde{R}$; complete $Q = [\tilde{Q}, Q']$ to a full unitary matrix, "padd" $\tilde{R}$ with 0 rows to get R, and once again $A = QR$. □

Corollary 1. *From the Gram-Schmidt process followed by a potential "pruning" of zero columns and rows, one can obtain the "condensed" QR factorization $A = \tilde{Q}\tilde{R}$, where $\tilde{Q} \in \mathcal{M}_{m,r}$ and $\tilde{R} \in \mathcal{M}_{r,n}$, the former having orthonormal columns, and the latter being upper triangular. The parameter r here can take the place of either n (if no pruning) or the rank of A (if pruning)*

Remark 1. *The following are easily seen to be true:*

- *We can choose R to have non-negative diagonal entries.*
- *If A is full rank, R can in fact be chosen to have positive diagonal entries. In this case the condensed QR factorization is* unique.
- *If A is full rank, the only non-uniqueness in the condensed QR form derives from the possibility to attach phases to the columns of $\tilde{Q}$. In other words, letting $D =$diag$(e^{i\theta_1}, \ldots, e^{i\theta_n})$ be a generic notation for diagonal matrices of phases, for any two factorizations $Q_1R_1 = Q_2R_2$ of A, there exists a D such that $Q_1 = Q_2D^*$ and $R_1 = DR_2$.*

Remark 2. *If we replace $\mathbb{C}$ with $\mathbb{R}$, then Q becomes orthogonal rather than unitary, and R is real. All the rest holds.*

One last thing that needs to be mentioned: Modified Gram-Schmidt (MGS) is relatively fast and stable. There are, however, more stable algorithms (and arguably better, computationally) for calculating QR; the two most important ones are the Householder reflector one (mentioned in the homework), and the Givens rotation one.

0.1 Using QR to solve Least Squares

A good part of classical numerical linear algebra (NLA) is concerned with solving the equation $Ax = b$ (... the joke amongst NLA people being that the rest of it is solving $Ax = \lambda x$.) Nevertheless, as we know, if $b \notin \mathcal{R}(A)$, no such solution exists. What can one do then? The answer is to change the question: rather than *solve*, focus on *approximate*.

The "Least Squares" approximation problem is to minimize $||Ax-b||_2$, or equivalently, $||b-Ax||_2$, for given A and b (naturally, if $Ax = b$ does have solutions, then the minimization problem reverts to finding a/the solution). Equivalently, and more simply, we choose to minimize the square of this quantity, namely $||Ax - b||_2^2$.

Generally speaking, the reason for choosing the 2-norm is because it has an associated inner product, which makes things a lot less complicated, by transforming the problem into a geometry one: finding the projection of b onto $\mathcal{R}(A)$. This is illustrated by the following (not very hard) theorem.

Theorem 1. *(The Projection Theorem, finite dimensional version) Let V be an inner-product space, and $S \in V$ be a finite dimensional subspace. Then*

1) $V = S \oplus S^\perp$; that is, $\forall z \in V$, $\exists! x \in S$, $y \in S^\perp$, such that $z = x + y$. Incidentally, $x = P_S z$, $y = P_{S^\perp} z = (I - P_S)z$.

2) Given $z \in V$, the x in 1) is the unique element of S which satisfies $z - y \in S^\perp$.

3) Given $z \in V$, the x in 1) is the unique element of V realizing the minimum $\min_{s\in S} ||z - s||_2^2$. The norm here is the one induced by the inner product.

Proof. The proof is in the picture.

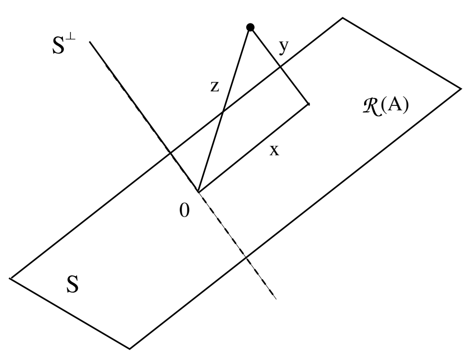

1) Let $\{\phi_1, \ldots, \phi_k\}$ be an orthonormal basis for S. Let $x = \sum_{i=1}^{k} \langle \phi_i, z\rangle \phi_i$, and let $y = z - x$. Then trivially $x \in S$, and

$$\langle \phi_i, y\rangle = \langle \phi_i, z\rangle - \langle \phi_i, x\rangle = 0 \ ,$$

so $y \in S^\perp$. Uniqueness for both x and y follows from the fact that $S \cap S^\perp = \{0\}$.

2) This is a restatement of 1).

3) Note that for $s \in S$, $z - s = x - s + y$, and since $x - s \in S$ and $y \in S^{\perp}$, we can write by the Pythagorean Theorem

$$||z - s||_2^2 = ||x - s||_2^2 + ||y||_2^2 ,$$

so clearly the LHS is minimized when $x = s$.

$\square$

0.1.1 Solving Least Squares with QR

We start with this, because it is the most widely used algorithm for solving Least Squares. Recall that we want to minimize $||b - Ax||_2^2$ for some matrix $A \in \mathcal{M}_{m,n}$ and vector $b \in \mathbb{C}^n$.

Assume wlog rank$(A) = n$, write $A = QR$, the full QR decomposition of A, and let $A = \tilde{Q}\tilde{R}$ be the condensed QR decomposition of A (so that $\tilde{R}$ has no zero rows, and $\tilde{Q}$ has orthonormal columns), and rearrange $Q = [\tilde{Q}, Q']$, $R = \begin{bmatrix} \tilde{R} \\ 0 \end{bmatrix}$.

(If A is not full rank, one can still do the condensed decomposition of A, find a "solution" x', and then "padd" it appropriately with zeros.)

Write

$$||b - Ax||_2^2 = ||QRx - b||_2^2 = ||Q^*(QRx - b)||_2^2 ,$$

since the 2-norm is invariant under unitary multiplication, and thus

$$||b - Ax||_2^2 = ||Rx - Q^*b||_2^2 = \left|\left| \begin{bmatrix} \tilde{R}x - \tilde{Q}^*b \\ (Q')^*b \end{bmatrix} \right|\right|_2^2 = ||\tilde{R} - \tilde{Q}^*b||_2^2 + ||(Q')^*b||_2^2 .$$

Thus the minimization problem is transformed into finding the minimum of $||\tilde{R}x - \tilde{Q}^*b||_2^2$, but, since we can choose $\tilde{Q}$ so that $\tilde{R}$ has all positive entries on the diagonal, $\tilde{R}$ is invertible, and $\tilde{R}x = \tilde{Q}^*b$ is solvable.

Thus, the minimum is obtained uniquely at the point for which $\tilde{R}x = \tilde{Q}^*b$.

Note that we only need matrix multiplication and computing QR to solve; we will see later why this is important.