Matrix Research PhL 12.14.04
finish 12.22.04

Now that I have a foundation of sorts from M&M and other work, I want to see what pieces are still missing, and I know there are many.

Question: What can you say about rank and the product of matrices? Here are a set of claims on web:

0 The rank of **A** is the dimension of its range. // OK
1 rank(**A**) = rank($\mathbf{A}^T$) = rank($\mathbf{A}^H$) // OK
2 rank(**A**) = maximum number of linearly independent columns (or rows). // OK
3 [**A**:*n*#*n*]: rank(**A**) + nullity(**A**) = *n* // OK
4 [**A**:*n*#*n*]: **A** is non singular iff rank(**A**)=*n*. // OK
5 [**A**:*m*#*n*, **B**:*n*#*k*]: rank(**A**) + rank(**B**) - *n* <= rank(**AB**) <= min(rank(**A**), rank(**B**)) // OK, see below!
6 rank(**A** + **B**) <= rank(**A**) + rank(**B**) // OK, see below!
7 [**X**:*m*#*m* and **Y**:*n*#*n* non-singular]: rank(**XA**) = rank(**AY**) = rank(**A**) // OK, see below!
8 rank(**KRON**(**A**,**B**)) = rank(**A**)rank(**B**) // OK, see below!
9 rank(**DIAG**(**A**,**B**,...,**Z**)) = sum(rank(**A**), rank(**B**), ..., rank(**Z**)) // OK, see below!

The last 5 items above are all new to me. Proofs of the others follow somewhere in the text here:
// all these items are now proven below.

Theorem M1: If C = AB, then r(C) ≤ min[r(A) , r(B)], A and B any conformable matrices.

Proof:
(1) If Bx=0, THEN Cx = 0 (not the other way), so x in N_B => in N_C so n(C) ≥ n(B). But since r+n= N, we have that r(C) ≤ r(B)
(2) Repeat the above analysis in Hermitian conjugate world where x'C' = xB'A' (' =HC), which is the adjoint * world of Stakgold. Here we conclude that r(C') ≤ r(A'). But of course r(X) = r(X'), so we conclude that r(C) ≤ r(A).
Combining these results gives: r(C) ≤ min [r(A), r(B)]

Theorem M2: If C = AB where A is non-singular, then r(C) = r(B) and vice versa A ↔ B.

Proof: From the previous theorem, we know that r(C) ≤ r(B) since r(A) = maximal. Since A is invertible, we can write B = A^{-1}C and conclude that r(B) ≤ r(C) by the same argument. Thus, r(B) = c(C). So multiplication by a non-singular matrix does not change rank.

Definition: **ERO** = Elementary Row Operation : Scale a row by nonzero, swap rows, or add multiple of a row. None of these things changes truth of the set of equations Ax = y.

Definition: **Row Echelon Form**.
This is what you get by applying ERO's to a matrix to clear out the lower triangle, the thing we called Gauss elimination in Scheid. Zero rows (if any) are put on the bottom. The leftmost entry in any row is a 1. The remaining entries in each row could be anything. You swap rows in an effort to maintain the leading diagonal of ones, but if you fail, then it jogs over.

Definition: **Reduced Row Echelon Form**. (aka **Normal Form**)
If you first go to the Row Echelon form and continue to clear out the upper triangle as well, this is what you get. I called this Gauss-Jordan elimination in my Scheid notes. In this form, each column has only a single entry. Happy columns have a sole 1, but where you had to jog over in doing Gauss elimination, you can have something other than a 1.

Definition: A is **Row Equivalent** to B means that you can get from A to B using EROs.

Definition: A is **Equivalent** to B means A is either row or column equivalent. Write **A.eq.B**. Notice that this is a different meaning from saying that

Theorem M3: The Reduced Row Echelon form of a matrix A is unique.

An induction proof is provided in the Lecture 17 notes here.
http://www.maths.usyd.edu.au:8000/u/don/courses/math1902/
For now I accept this, it seems awefully reasonable, since there is a unique sequence of operations you do to get to the reduced echelon form!

Theorem M4: EROs do not change the rank of a matrix.

Proof: Obviously non-zero scaling or row swapping don't change the rank. Doing $r_2 = r_2 + \alpha r_3$ can be put into the det formula, and you find that you are just adding 0 to the det because the second piece has two rows the same. This applies to all determinants that include segments of row r_2. This, this operation does not change any determinants at all, so the rank cannot change.

Theorem M5: EROs do change the determinant of a matrix in a predictable manner, but cannot cause a non-vanishing determinant to vanish.

Proof: Swapping rows makes a - sign. Scaling a row by $\alpha \neq 0$ scales the det by α. Doing the addition as above has no effect on det. The final fact follows.

Theorem M6: An ERO is invertible and thus non-singular.

Proof: For scaling we know invertible, and same for row swap. For $r_2 = r_2 + \alpha r_3$, the inverse does $r_2 = r_2 - \alpha r_3$. Thus, each basic type of ERO is invertible, so any product of them is.

Theorem M7: The rank of a matrix equals the number of non-zero rows in its reduced row echelon form.

Proof: Each non-zero row has a 1 at its left end, so to speak. From the point of view of doing a determinant, we can swap columns to get these 1's all lining up on the diagonal. The det of the upper-left matrix is then non-zero, since it is diag(1,1,1..) for as many rows as are non-zero. Pretty spiffy.

Theorem M8. Row and Column Operations on A can be done by A' = RAC, and R and C are non-singular.

Proof: Consider this where r1 is the first row of the matrix A, and the left matrix is R.

$$\begin{pmatrix} 2 & 3 & 4 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}\begin{pmatrix} r1 \\ r2 \\ r3 \end{pmatrix} = \begin{pmatrix} 2*r1+3*r2+4*r3 \\ r2 \\ r3 \end{pmatrix}$$

It is clear that we can do any elem row op we want on row 1 in this way, and indeed on any row, QED. A similar argument can be made in terms of columns with C acting from the right. As for the final statement, consider just A' = RA. Since R cannot make the det vanish, we must have $\det R \neq 0$.

Examples of EROs.

1. Suppose you want to clean out the first column of a matrix below the top element. As long as the leftmost element in row r1 does not vanish, you can do it like so: (here, **r1** is a row vector)

$$\begin{pmatrix} 1 & 0 & 0 \\ \alpha & 1 & 0 \\ \beta & 0 & 1 \end{pmatrix}\begin{pmatrix} \mathbf{r1} \\ \mathbf{r2} \\ \mathbf{r3} \end{pmatrix} = \begin{pmatrix} \mathbf{r1} \\ \mathbf{r2}+\alpha*\mathbf{r1} \\ \mathbf{r3}+\beta*\mathbf{r1} \end{pmatrix} = \begin{pmatrix} \mathbf{r1} \\ \mathbf{r2'} \\ \mathbf{r3'} \end{pmatrix} \qquad E_1 R = R' \qquad \text{// R is some general 3x3 matrix}$$

2. Now suppose you want to clean out the second column below the diagonal. As long as the second element of r2' does not vanish, you can do it like so:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \gamma & 1 \end{pmatrix}\begin{pmatrix} \mathbf{r1} \\ \mathbf{r2'} \\ \mathbf{r3'} \end{pmatrix} = \begin{pmatrix} \mathbf{r1} \\ \mathbf{r2'} \\ \mathbf{r3'}+\gamma*\mathbf{r2'} \end{pmatrix} = \begin{pmatrix} \mathbf{r1} \\ \mathbf{r2'} \\ \mathbf{r3''} \end{pmatrix} \qquad E_2 R' = R''$$

3. **Theorem M9.** (**The LU Factorization**.) Assuming there are no zero elements in the wrong place, you can do this process for each of the columns in a square matrix until it is in UT form. Since all the ERO matrices we have used (adding a scaled row to a lower row) are LT, *their product* Π is LT (I have proven this fact elsewhere: UT*UT = UT and same for LT). Moreover, the diagonal elements of P will be all ones, which you can see by just doing it for two LT's with unit diagonal elements. Thus, detP = 1 since it is LT and has all 1's on the diagonal. So, we have

$$\Pi * R = UT$$

Now let $L = \Pi^{-1}$, and R = general matrix A, and UT = U, we then have

$$\Pi A = U \quad \Rightarrow \quad A = LU \qquad L = \Pi^{-1} = (E_1 E_2 \ldots)^{-1} = \ldots E_2^{-1} E_1^{-1}$$

Now L will also have 1's on the diagonal which you can see looking at $\Pi L = E$.

So if you simply do the above Gauss elimination process and keep track of the ERO matrices and then you know L and U. Another way is to just do the Gauss to get U, then compute $L = AU^{-1}$.

4. Note that the row swap operation is NOT lower triangular, so we if we have to do any of these in our ERO processing above, we do not get the A = LU form. Here is a swap of the first two rows, and it is clearly neither LT nor UT. It is symmetric, however. This is a "permutation" matrix.

$$\begin{pmatrix} 010 \\ 100 \\ 001 \end{pmatrix} \begin{pmatrix} \mathbf{r1} \\ \mathbf{r2} \\ \mathbf{r3} \end{pmatrix} = \begin{pmatrix} \mathbf{r2} \\ \mathbf{r1} \\ \mathbf{r3} \end{pmatrix}$$

5. One could adjust all the row positions with a set of swaps to make the Gauss operation work. These swaps are like the one in 4 above. The combination of all of them we shall call P^T so our results from 3 are then

$\Pi\ (P^T)A = U \Rightarrow (P^T)A = LU \Rightarrow A = PLU$

Theorem M10: (**PLU Factorization**). Any non-singular square matrix A can be written as PLU where P is a permutation matrix, L is lower-triangular with 1's on the diagonal, and U is upper triangular. The matrix P is a some permutation of the unit column vectors. [Maple does this with LUdecomp]

Definition: A and B are **equivalent** if there exist non-singular P and Q where A = PBQ.

Theorem M11: Equivalent matrices have the same rank.

Proof: Assume A = PBQ with P and Q non-singular. We know that r(BQ) = r(B) from our earlier theorem, and then r(A) = r(BQ) by the same argument. Thus, r(A) = r(B). QED.

Theorem M12: (**Submatrix Theorem**) Let matrix B be a horizontal slice of NxN matrix A, where B is S rows high. Then $r(A) \geq r(B) \geq r(A) + S - N$. .

Proof: The first part is obvious that $r(B) \leq r(A)$ if B is only a piece of A. The other inequality is proved as follows (see old notes!). (1) slide B to the top, then write it in normal form. ranks are unchanged. (2) Slide its zero rows to the bottom if there are any. (3) Then take the resulting upper part to normal form. The resulting matrix still has the rank of A, but has S - r(B) rows of zeros on the bottom. The rank of the resulting matrix must be ≤ the number of upper rows, which is N - (S - r(B)). So $r(A) \leq N - S + r(B)$. This is the desired result. It is all very obvious if you draw pictures.

Theorem M13: (**Sylvester's Law of Nullity**) If A and B are NxN, then $r(AB) \geq r(A) + r(B) - N$. In terms of nullity, $n(AB) \leq n(A) + n(B)$. Our previous theorem said $r(AB) \leq \min[r(A),r(B)]$ so the combined result is then

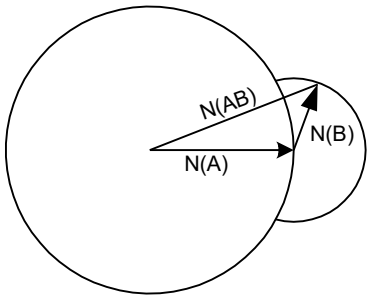
$$r(A) + r(B) - N \leq r(AB) \leq \min[r(A),r(B)]$$

Rewrite as

$$\max[n(A),n(B)] \leq n(AB) \leq n(A) + n(B) \qquad \text{// triangle rule!}$$

Proof: Write $XY = (PX_{nf}Q)Y$ where X is in normal form and P and Q got A to be in normal form. Then we have $XY = P(X_{nf}QY)$. But the matrix $X_{nf}QY$ is a strip of the matrix QY of height S = r(X), so we apply our submatrix theorem $r(B) \geq r(A) + S - N$ with A = QY and B = $X_{nf}QY$ to get $r(X_{nf}QY) \geq r(QY) + r(X) - N$ which we rewrite as $r(XY) \geq r(Y) + r(X) - N$. QED. The nullity part comes from the usual r(A) + n(A) = N relationship. [Again, see hand written notes for help if needed.]

Interpretation of the Triangle Rule: $\max[n(A),n(B)] \le n(AB) \le n(A) + n(B) n(AB) \le n(A) + n(B)$



Well, there is a triangle, but there is more to it that just the triangle. We assume N(A) > N(B) for this picture. Then the top of the imagined vector N(B) can only be on the portion of the small circle that is shown, because the N(AB) distance must stay larger than N(A). All this rule says is that the N(AB) is larger than either one, but cannot exceed the sum of the two.

Theorem M14: $r(A+B) \le r(A) + r(B)$

Proof: Consider rank as the dim(Range). Consider $(A+B)x = y$. Define $y_1 = Ax$ and $y_2 = Bx$. In worst case, these ranges might not overlap at all in terms of basis vectors, so $\dim(R_{A+B})$ = sum of range dims. that is, the vector $y = y_1 + y_2$ is then spanned by the set including all basis vectors ion each set. In all other cases, there could be some overlap of the basis vectors. In worse case, there could be 100% overlap. Then we would have $r(A+B) = r(A) = r(B)$. Both extremes are compatible with $r(A+B) \le r(A) + r(B)$.

Definition: A **Householder Matrix**: given any real vector u , the matrix is $H = 1 - \sigma uu^T$ where the number $\sigma = (2/u^Tu)$. Think of H as H(u). [aka an **Elementary Reflector**] Can write as $H = 1 - 2ww^T$ if use w that is normalized.

Theorem 15: The following are true for a Householder Matrix H:
(1) $H = H^T$ which is true for any constant σ
(2) $HH = 1$ (or $H^{-1} = H$) which is true only for the specific σ shown in the definition
(3) $H^T = H^{-1}$ so H is real orthogonal (follows from 1 and 2)
(4) $(uu^T)^2 = (u^Tu)(uu^T)$ is true for any u
(5) $(uu^T)v = (u^Tv)u$ is true for any u and any v

Proof: Item 1 is trivial. Items 4 and 5 are trivial if you just write them out. Use them to show item 2. This is all done very nicely in Jerry Shultz notes which I have printed and marked up.

Theorem M16: The product of a Householder matrix and a real orthogonal matrix is a real orthogonal matrix.

Proof: Let $R = H_1S$ with S = ro. Then $R^{-1} = (H_1S)^{-1} = S^{-1} H_1^{-1} = S^TH_1=(H_1S)^T = R^T$.

Corollary: The product of two Householder matrices is a real orthogonal matrix.

Proof: A Householder matrix is real orthogonal, so the theorem shows it to be true.

Theorem M17: One can find a Householder matrix H such that Hv = (β,0,0...) for any vector v, where the value of β must be $\beta = \pm \text{sqrt}(v_1^2 + v_2^2 + ... + v_N^2) = \pm |v|$. Usually one takes β to have the same sign as v_1.

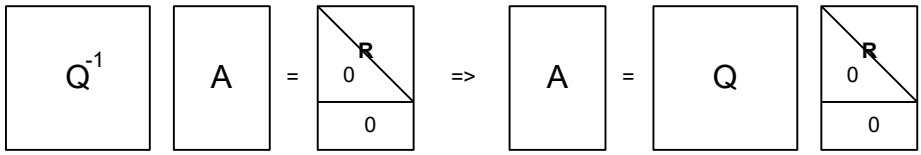
Proof: If you select u = $(\alpha, v_2, v_3...)$ where $\alpha = v_1 \pm \text{sqrt}(v_1^2 + v_2^2 + ... + v_N^2)$, and just write things out, you find that the result is true. One fact along the way is that you need $2u^Tv = u^Tu$. This is all clearly shown in Jerry Shultz's notes.

Theorem M18: (**QR Factorization**). Any square matrix A can be written as A = QR where Q is real orthogonal, and R is upper triangular.

Proof: Consider the first column of matrix A as vector v and find H_1 which brings it to (β,0,0...) form. We know we can do this by the previous theorem. Then to the resulting matrix H_1A with this cleaned up first column, apply the matrix $H_2 = \text{diag}(1, H_2')$ where H_2' is N-1 x N-1 and acts on the second column's lower N-1 components to clean out the bottom of that column, so it has the form (a,β',0,0,...). Then keep doing this, reducing the dimension each time, until you have a result that is upper triangular! You then have that $H_N H_{N-1} .....H_2H_1A = R$ where R is the upper triangular result. The product of the H's is real orthogonal according to previous theorem. Call it Q^{-1} to get $Q^{-1}A = R$, then left multiply by Q, A = QR.

Theorem M19: (QR factorization for non-square matrices) $A = Q \begin{pmatrix} R \\ 0 \end{pmatrix}$ or $A = Q\,(R_1, R_2)$ as sketched below.

Proof: If A is "tall", N x n, we just apply the previous theorem to start clearing out the columns of A, using "large" H matrices, which multiplied together give Q^{-1}, like so:

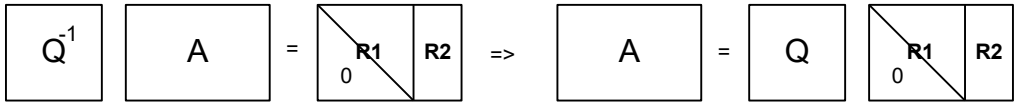


When we get done, the bottom N-n rows of A' will be all zeros below its triangular region. So write this as

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix}$$

where Q is the large NxN, and R is the small nxn. You can see that the rightmost columns of Q are doing nothing here since they hit those bottom rows of zeros, so write $A = (Q_1, Q_2)$ and then $A = Q_1R$ and everything is n x n.

On the other hand, if A is "wide", n X N, then we have this picture



and then we can write A = Q (R_1, R_2) where R_2 is completely general.

These notations are used at this site: http://www.netlib.org/lapack/lug/node40.html

Theorem M20: Gauss Elimination gives a different QR factorization where Q is non-singular but not necessarily real-orthogonal. So the QR is a sort of messy normalized Gauss elimination. BUT, whereas Gauss sometimes has to "skip a column" along the diagonal edge, the QR never does this! The QR has a reliable diagonal boundary with no zeros along it. [I think this is not true!]

In Gauss elimination, we apply ERO's to reach upper triangular form. From our theorem above, you can think of each such ERO as A' = RA where R is non-singular. If you call the final triangular form T, you end up with T = XA where X is the product of all the EROs. Since we know this product is invertible, call it $X = Q^{-1}$. We then have $T = Q^{-1}A$ which means A = QT. Renaming Q to be R, we get A = QR. Notice that in this factorization, all we know is that Q is non-singular, it is not real-orthogonal.

Theorem M21: rank(A⊗B) = rank(A)rank(B)

Proof: Write A = Qa and B = qb to get A⊗B = Qa⊗qb = (Q⊗q)(a⊗b) using two QR factorizations. Since Q and q are RO they have full rank, so rank(A⊗B) = rank(a⊗b), see theorem earlier that r(AB) = r(B) if A is full rank. Visualize a⊗b as triangle cross triangle where matrix b appears as a whole, elements like a_{11}b. Each fat row of this thing has r_b nonzero fine rows, and there are r_a of these fat rows that are non-zero, so number of non-zero fine rows is then $r_a r_b$. Thus, rank(A⊗B) = r(A)r(B). So this proof required the QR factorization1

Theorem M22: rank(A⊕B) = rank(A) + rank(B)

Proof: Start with A⊕B = diag(A,B). First apply diag(Q,1) to convert A to triangular form QA, then apply matrix diag (1,q) to do the same to B. Then you have diag(A_t, B_t) . The total number of non-zero rows in this matrix is $r_A + r_B$, QED. This is a direct sum thing.

****

Generalization to Complex Householder

Definition: A **Householder Matrix**: given any complex vector u , the matrix is $H = 1 - \sigma uu^H$ where the real number $\sigma = (2/u^Hu)$. Think of H as H(u). [aka an **Elementary Reflector**] Can write as
$H = 1 - 2ww^H$ if use w that is normalized.

Theorem M23: The following are true for a Householder Matrix H: (superscript means T* = Herm adjoint)

(1) $H = H^H$ which is true for any constant σ

(2) HH = 1 (or $H^{-1} = H$) which is true only for the specific σ shown in the definition

(3) $H^H = H^{-1}$ so H is unitary (follows from 1 and 2)

(4) $(uu^H)^2 = (u^Hu)(uu^H)$ is true for any u

(5) $(uu^H)v = (u^Hv)u$ is true for any u and any v

Theorem M24: The product of a Householder matrix and a unitary matrix is a unitary.

Proof: Let $R = H_1S$ with S = uni. Then $R^{-1} = (H_1S)^{-1} = S^{-1} H_1^{-1} = S^H H_1 = (H_1S)^H = R^H$.

Corollary: The product of two Householder matrices is a unitary matrix.

Proof: A Householder matrix is unitary, so the theorem shows it to be true.

Theorem M25: One can find a Householder matrix H such that $Hv = (\beta,0,0...)$ for any vector v, where the value of β must be $\beta = \pm \text{sqrt}(|v_1|^2 + |v_2^2| + ... + |v_N^2|)e^{i\phi} = \pm |v|e^{i\phi}$ where ϕ is the phase of v_1.

Proof: We select $u = (\alpha,v_2,v_3...)$ as before, and $\sigma = 2/(u^Hu)$ real as before. Out condition as before is that $u^Hu = 2u^Hv$. This becomes $|\alpha|^2 + S = 2[\ \alpha^* v_1 + S\ ]$ so end up now with $|\alpha|^2 - 2v_1\alpha^* - S = 0$, where $S = \Sigma |v_\ell|^2$ = real, excluding the first term. Compare this to Jerry's equation top of page 3. Remember that we get $\beta = (Hv)_1 = v_1 - u_1 = v_1 - \alpha$. I find it easiest to say $\alpha = |\alpha|\ e^{i\theta}$ and $v_1 = |v_1|e^{i\phi}$ and the resulting equations are

$$|\alpha|^2 - S - 2\ |\alpha|\ |v_1| \cos(\phi-\theta) = 0 \quad \text{and} \quad -2\ |\alpha|\ |v_1| \sin(\phi-\theta) = 0$$

If we let $\phi - \theta = n\pi$, the second is happy, pick n = 0, then first says $|\alpha|^2 - S - 2\ |\alpha|\ |v_1| = 0$. Solving this for $|\alpha|$ gives $|\alpha| = +\ |v_1| \pm |v|^2$, but second term is larger in general, so assume $|\alpha| = |v_1| + |v|^2$. This then gives $\beta = v_1 - \alpha = |v_1|e^{i\phi} - [|v_1| + |v|^2]\ e^{i\phi} = -\ |v|^2e^{i\phi}$. If we chose instead n = 1, then $|\alpha|^2 - S + 2\ |\alpha|\ |v_1| = 0$ and then $|\alpha| = -\ |v_1| \pm |v|^2$ and again must take second plus so $|\alpha| = -\ |v_1| + |v|^2$ and $\beta = +\ |v|^2e^{i\phi}$. So the general conclusion is that $\beta = \pm\ |v|^2e^{i\phi}$ when ϕ is the phase of v_1 . So if v_1 is complex, so is β.

Theorem M26: (**complex QR Factorization**). Any square matrix A can be written as A = QR where Q is unitary, and R is upper triangular.

****

Type II Digression

This site: http://www.netlib.org/lapack/lug/node128.html#secorthog suggests a way to make β be real. I have tried to pursue this a bit here.

Definition: A **Type II Householder Matrix**: given any complex vector u , the matrix is $H = 1 - \sigma\ uu^H$ where σ is a complex number satisfying $|\sigma|^2\ u^Hu = 2\text{Re}(\sigma)$.

Theorem M27: The following are true for a Type II Householder Matrix H:
(1) $H \neq H^H$
(2) $H^*H = 1$ (or $H^{-1} = H^*$) which is true only for the specific σ shown in the definition
(3) H is not unitary
(4) $(uu^H)^2 = (u^Hu)\ (uu^H)$ is true for any u
(5) $(uu^H)v = (u^Hv)u$ is true for any u and any v

The claim is that this gives β real, but the product of the H's won't be unitary, which seems not good, so let's let this dead dog lie for now.

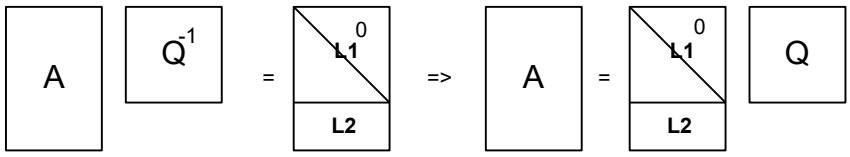
****

Theorem M28: (**complex LQ Factorization**). Any square matrix A can be written as A = LQ where Q is unitary, and L is lower triangular.

Proof: The proof is pretty much the same as for the QR except: (1) We apply our H matrices from the right instead of the left, and in doing so (2) they clear out rows instead of columns, which is why you end up with a lower triangular instead of an upper triangular. You get A' = AHHHH and write the product of the H's as Q^{-1} which is unitary, so then A' = L = AQ^{-1}. Now right multiply by Q to get LQ = A and you have arrived!

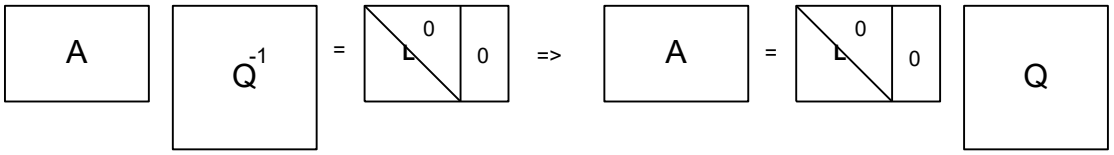
Theorem M29: (LQ factorization for non-square matrices) A = $\begin{pmatrix} L1 \\ L2 \end{pmatrix}$ Q or A = LQ_1 as sketched below.

Proof: If A is "tall", N x n, we just apply the previous theorem to start clearing out the rows of A, using "small" H matrices, which multiplied together give Q^{-1}, like so:



If we write L = $\begin{pmatrix} L1 \\ L2 \end{pmatrix}$ then this says A = $\begin{pmatrix} L1 \\ L2 \end{pmatrix}$ Q. Or just A = LQ.

If A is "wide: n x N, then we get this picture



and now the last rows of Q do nothing. If we write Q = $\begin{pmatrix} Q1 \\ Q2 \end{pmatrix}$ then we get A = LQ_1.

Theorem M30: The inverse T^{-1} of an upper triangular matrix T is upper triangular, and the diagonal elements of T^{-1} are the inverses of those of T. This only makes sense if all diagonal elements of T are non-zero, otherwise detT = 0 and T^{-1} does not exist.

Proof: This is an odd theorem. I first tried to prove it by showing why cofactors vanish, but the thing was not simple. I then found a proof here: www.ualberta.ca/~kumar/handouts/chap3.pdf where you just show the

idea that T B = 1 and you do things one at a time. I think this is the same as my proof which follows. Consider the method of Scheid, page 54 of my Scheid notes. You start off with a table like this:

*	*	*	1	0	0
0	*	*	0	1	0
0	0	*	0	0	1

where I am just showing a 3x3 example. I is the rightmost 3 columns. The idea is that if you can just do ERO's and get the left thing to be I, then the right thing is the inverse. I called this Matrix Inversion by Gauss-Jordan Elimination in those notes. We know that our T matrix on the left has all non-zero elements on the diagonal, so we can just scale each row separately to get ones.

1	*	*	*	0	0
0	1	*	**0**	*	0
0	0	1	**0**	**0**	*

Now we just do our row additions to clear out the *'s in the left 3 columns. We always add lower rows to upper rows to do this. The reason for this is that we don't have to clear out anything in the bottom left half because it is already cleared since T is triangular. In doing such additions, there is no way to destroy three bolded zeros. That is the key. When you are done, you have this

1	0	0	*	*	*
0	1	0	**0**	*	*
0	0	1	**0**	**0**	*

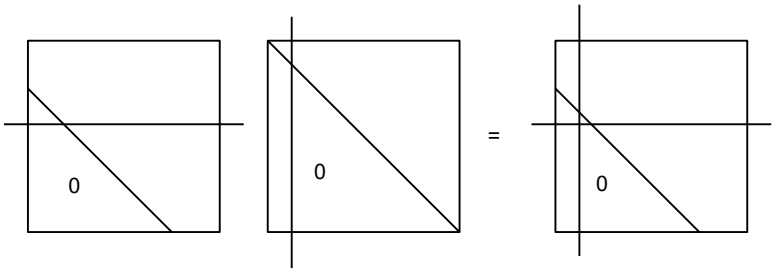
and the result is that T^{-1} is also upper triangular! The diagonal elements are the inverses of those of T.

Theorem M31: The product of two upper triangular matrices is upper triangular.

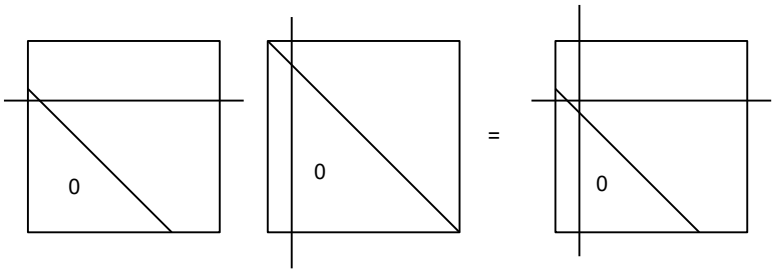
Proof: Just draw it on paper and the result is obvious.

Theorem M32: Let A = triangular, B = partially triangular in same sense. Then C = AB and D = BA are both partially triangular in the same sense as B.

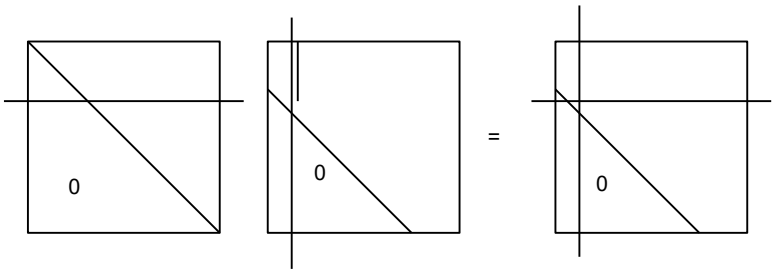
Proof: Here is a picture showing the idea that BA = D



For the matrix element selected by the crosshairs, you get zero because there are enough zeros in the row of B to carry us through into the zeros of the column of A. If we look a little higher,



now the zeros of the row of B are not enough and we get nonzero. If matrices are in the other direction, we get pictures like this for AB = C



and the conclusions are the same. In this specific case, I copied the segment on the left and rotated it 90 degrees to show that it is shorter than the vertical segment, so the result is nonzero.

Formal proofs of these things can be done with $A_{rc} = A_{rc}\ \theta(\ r \le c)$ Heaviside, for an upper triangular matrix. Then you could have $B_{rc} = B_{rc}\ \theta(\ r \le c|{+}m)$, where m would be 1 for upper Hessenberg. Here is a proof for the case C = AB :

$$C_{ij} \ = \Sigma_\ell \ A_{i\ell}\ \theta(\ i \le \ell\ )\ B_{\ell j}\ \theta(\ell \le j + m\ )$$

Just looking at the insides of the two theta functions shows that we need $i \le \ell \le j + m$, which implies that $i \le j + m$, otherwise there are no terms in the sum! But this says that $C_{ij} = C_{ij}\ \theta(i \le j + m)$ which means it has the same triangular shape as B. If we swap the location of the "m", we get $i \le \ell + m$ and $\ell \le j$, and this says that $i - m \le \ell \le j$ so we need $i - m \le j$ which again says $i \le j + m$ and we get the same conclusion.

We could further generalize this theorem to say that the result of two triangular's of the same sense has the sense of the triangular with the least zeros.

Theorem M33: If you multiply two (partially) triangular matrices of the same sense, the result has the shape of the factor matrix having the fewest zeros, ie, the one that is least triangular.

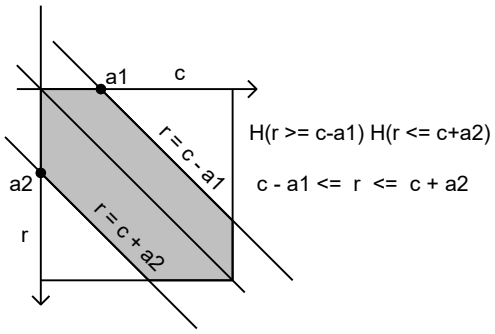
Theorem M34: In the QR iteration scheme we factor A = QR, where Q is real orthogonal and R is upper triangular, and then define $A' = RQ = Q^{-1}\ A\ Q = Q^T\ A\ Q$. The transformation from A to A' preserves Hessenberg form.

Proof: Since R is triangular, so is R^{-1}. We then have $Q = AR^{-1}$. If A is upper Hessenberg (meaning partially triangular), then so is Q by our last theorem or by the previous one. Looking then at A' = RQ, we see that again we have triangular times Hessenberg, so the result is again Hessenberg. QED.

Corollary: The same transformation as above preserves tridiagonality. WRONG!

Proof: A triadiagonal matrix is both upper Hessenberg and lower Hessenberg at the same time. Thus, the transformed matrix will be both these things, and must therefore be tridiagonal. (Is this really true?? is there not a sense conflict?) ARGUMENT HERE IS WRONG!

Definition. A **banded matrix** A with $c - a_1 \le r \le c + a_2$ means that $A_{rc} = 0$ for r,c outside the range shown. Here is a picture of this banded matrix A (integers a_1 and a_2 are assumed non-negative)



Example #1: An **upper triangular** (**UT**) matrix A has a1 = N and a2 = 0

Example #2: An **generalized upper Hessenberg** (**UHn**) matrix has a1 = N and a2 = n. A standard issue upper Hessenberg has n=1 and we write as UH1.

Example #3: A **tridiagonal** matrix has a1 = a2 = 1.

Theorem M35: If A is a banded matrix with $c - a_1 \le r \le c + a_2$, then if $a_1 \ge N$ we can ignore the left inequality, and if $a_2 \ge N$ we can ignore the right inequality.

Proof: If $a_1 \ge N$, the upper banding boundary in the above picture is above the entire matrix and so puts no restrictions on the matrix non-zero area. The entire UT area is non-zero. Saying this another way, if $a_1 \ge N$ then write $a_1 = N + p$ where $p \ge 0$. Then we have $r \ge c - a_1$ or $r \ge (c-N) - p$. At either extreme of the range of c, we get $r \ge -p$ and $r \ge -N -p$, neither of which restricts anything since $r \ge 0$. A similar argument can be made for $a_2 \ge N$, see the above picture.

Theorem M36 (Banded Matrix Theorem): Consider C = AB where A is banded by $c - a_1 \le r \le c + a_2$ and where B is banded by $c - b_1 \le r \le c + b_2$. Then C is banded by $c - c_1 \le r \le c + c_2$, where $c_1 = a_1 + b_1$, $c_2 = a_2 + b_2$.

Proof: Write $C_{ij} = \Sigma_\ell A_{i\ell} B_{\ell j}$ and then translate the two input bands to

$$\ell - a_1 \le i \le \ell + a_2 \qquad\qquad j - b_1 \le \ell \le j + b_2$$

Restate the first pair with ℓ in the center to get this set, leave second pair as is:

$$i - a_2 \le \ell \le i + a_1 \qquad\qquad j - b_1 \le \ell \le j + b_2$$

As written, each pair separately contains very little info. Left pair says, for example, $i - a_2 \le i + a_1$ in order to have contributing terms in the sum, but this just says $- a_2 \le a_1$ which is true by assumption. But now exchange the two right sides of the above pair of pairs to get,

$i - a_2 \le \ell \le j + b_2$ $\quad$ $j - b_1 \le \ell \le i + a_1$

which we rewrite as

$i - a_2 \le j + b_2$ $\quad$ $j - b_1 \le i + a_1$

and again as

$i - j \le (a_2 + b_2)$ $\quad$ $i - j \ge - (a_1 + b_1)$

Now replace i,j with the clearer r,c indices to get

$r - c \le (a_2 + b_2)$ $\quad$ $r - c \ge - (a_1 + b_1)$

and these become

$r \le c + (a_2 + b_2)$ $\quad$ $r \ge c - (a_1 + b_1)$

and finally

$c - (a_1 + b_1) \le r \le c + (a_2 + b_2)$

which describes the product C as a banded matrix as claimed, QED.

Corollary 1. $UH_n * UH_m = UH_{n+m}$

Proof: For $A = UH_n$ we have $a_1 = N$ and $a_2 = n$. For $B = UH_m$ we have $b_1 = N$ and $b_2 = m$. Therefore, for $C = AB$ we have $c_1 = 2N$ and $c_2 = n+m$, so only condition is $r \le n+m$ which is UH_{n+m}

Corollary 2: $UH_0 * UH_m = UH_m$ or $UT*UH_m = UH_m$

Corollary 3: $UH_0 * UH_0 = UH_0$ or $UT*UT = UT$

Corollary 4. $UT * Tridiag = UH_1$

Proof: For A we have $a_1 = N$ and $a_2 = 0$. For B we have $b_1 = 1$ and $b_2 = 1$. Then applying the rule that $c - (a_1 + b_1) \le r \le c + (a_2 + b_2)$ tells us that $c - (N+1) \le r \le c + 1$ or simply $r \le c + 1$ which is UH1.

Corollary 5. Let C = AB where A is general and B is tridiagonal. Then C is general.

Proof: For A we have $a_1 = N$ and $a_2 = N$. For B we have $b_1 = 1$ and $b_2 = 1$. Then applying the rule that

$c - (a_1 + b_1) \leq r \leq c + (a_2 + b_2)$ tells us that $c - (N+1) \leq r \leq c + (N+1)$ which is general.

Theorem M37: The QR Transformation preserves UHn.

Proof: We start with the A = QR factorization of general NxN matrix A. Since R is UT, so is R^{-1} by theorem above. Write $Q = AR^{-1}$ = UHn * UT = UHn by our Corollary 2. Then A' = RQ = UT*UNn = UHn by the same corollary, so we have shown that A = UHn => A' = UHn. QED.

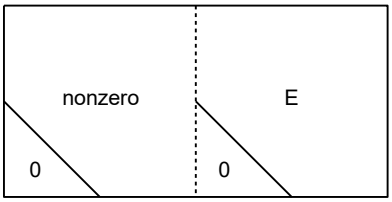
Theorem M38: The QR Transformation preserves tridiagonality if A is symmetric

Proof: We start with the A = QR factorization of general NxN matrix A. Since R is UT, so is R^{-1} by theorem above. Write $Q = AR^{-1}$ = tri * UT = UH1 by our Corollary 4. Then A' = RQ = UT*UH1 = UH1 by the same corollary. But if A is symmetric, we know that A' is as well, and in that case, UH1 => tridiagonal.

****

Theorem M39: The inverse of a UHn matrix is UHn.

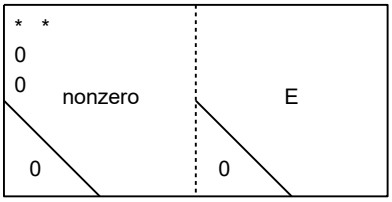
Consider the setup for doing Gauss-Jordan inversion



First column:

(a) There must be some non-zero element in the first column above row R where the triangle of zeros starts, otherwise detA = 0 and we are not invertible. If the top left element is 0, then do a swap to get a non-zero element into top left. Such a swap cannot affect any zeros in the lower triangles.

(b) Now add multiples of row 1 to rows 2 through R-1 as needed to clear out the first column. Such additions cannot affect the zeros in the triangles because they are below row R-1. After this point, we will never use row 1 again, so no later additions can ever change the first column of 0's in the triangular areas. We now have:



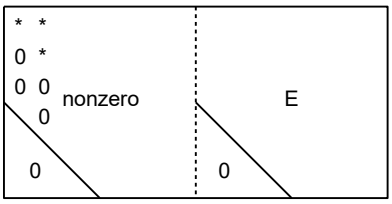
Second column:

(a) If the second column below the first element is all zeros, we don't have to clear it. Otherwise, there is some non-zero element between rows 2 and R+1. If necessary, we do a swap to get a non-zero element

as the second element in column 2. Such a swap only swaps two 0's in the left column (of each of our squares), and so cannot alter the triangular areas.

(b) Now add multiples of row 2 to rows 3 through R as needed to clear out the first column. Such an addition cannot change the triangle of zeros. After this point, we will never use row 2 again, so no later additions can ever change the second column of 0's in the triangular areas.

We now have:

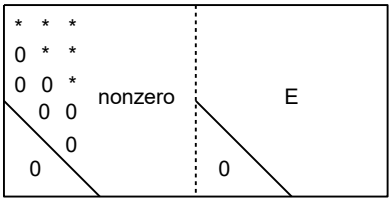


Third column:

(a) If the third column below the second element is all zeros, we don't have to clear it. Otherwise, there is some non-zero element between rows 3 and R+2. If necessary, we do a swap to get a non-zero element as the 3rd element in column 32. Such a swap only swaps 00 with 00 (of each of our squares), and so cannot alter the triangular areas.

(b) Now add multiples of row 3 to rows 4 through R+2 as needed to clear out the first column. Such an addition cannot change the triangle of zeros. After this point, we will never use row 3 again, so no later additions can ever change the third column of 0's in the triangular areas.

We now have:



We continue on in this way. For each column, we might have to do a row swap to get a non-zero element up to the pivot point, but such a swap does not affect the lower triangle of zeros.

When we are finished with this phase of the process, the matrix on the left is upper triangular with non-zero diagonal elements. We then begin the "Jordan" part of our clearing out process. We add multiples of the first column to columns to the right and we clear out the first row, and we do not affect our triangular areas. In this phase, we never have to do any column swaps because our pivot is a non-zero diagonal element. Thus we only add rows to the right, and that can never change the lower triangle.

After doing this, the left square is completely diagonal and our zeros are preserved in both triangles. We then just rescale each row to be 1 on the diagonal of the left square which is then E, and the right square contains the inverse. Thus, the inverse is UHn, same shape as the original matrix on the left.

Comments: Notice that this proof does not work if we use a doubly-banded matrix to start with. The failure of this proof is shown in the next section.

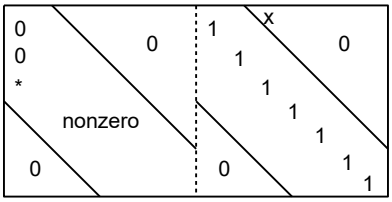
Corollary: The inverse of a UT matrix is a UT matrix.

Proof: This follows because UT = HM0 and the theorem is true for all HMn.

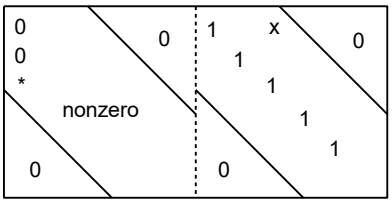
****

Hypothetical Theorem: The inverse of a banded matrix is a banded matrix of the same shape???

Not true: Proof: Consider the setup for doing Gauss-Jordan inversion

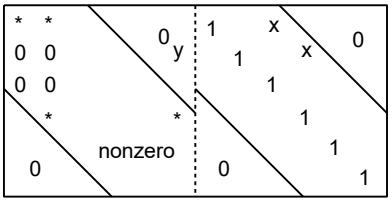


and assume a worst-case situation where we are going to need to swap row 1 with row R-1 to get our pivot to the top left corner. When we do this swap, a "1" in the right swaps up to where the x is, and pollutes our upper right triangle. This won't happen, however, if the upper right triangle is the not larger than the lower triangle. So let's make the two triangles the same size, (and we have 6 1's for illustration).



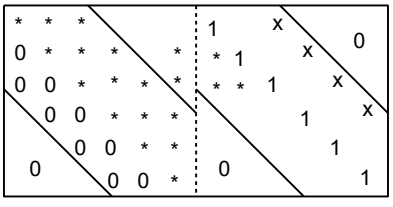
Now the upper triangle in the right square stays clean because the x is outside of it now.

Now let's assume a similar worst-case situation in the second column,



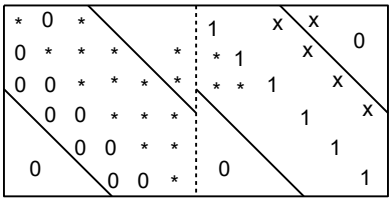
If we now swap the 4th row with the 2nd row to get our pivot in column 2, this creates the second x shown on the right. We are still OK there. However, the top triangle in the left square just got polluted where the y now sits, but we don't care about this triangle.

If we keep going in this manner clearing out the columns of the left square below the diagonal, in the right square both the top and bottom triangles will be untouched! We how have this situation:

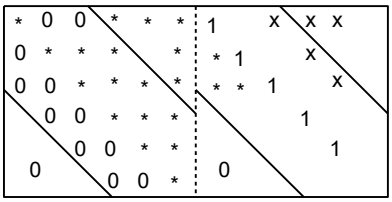


where we don't show all detail on the right, but triangles are clear.

Now we move to the problem of clearing out the rows, the "Jordan" phase. All diagonals in the left square are non-zero, they were our pivots for the Gauss phase. So we add a multiple of row 2 to row 1 to clear out one element to get



and BANG, we just killed our theorem! The reason is that in doing this addition, we caused something to appear in the upper right triangle! But let's keep going. We now add a multiple of the 3rd row to the 1st row to clear out the next element in the first row, and we take another hit!



Now the upper triangle in the left has become more polluted, and we have to clear out more elements of the first row, and this further pollutes the upper triangle on the right, and eventually it gets all filled in.
So our theorem is NOT TRUE!

****

Hypothetical Theorem: The inverse of a tridiagonal matrix is a tridiagonal matrix ???

Trying the above proof shows this is **not true**, and it is easy to show a counterexample,

S := matrix([[1, 2, 0], [4, 5, 6], [0, 8, 9]])

inverse = matrix([[1/25, 6/25, -4/25], [12/25, -3/25, 2/25], [-32/75, 8/75, 1/25]])

****

Addendum

Theorem M40: If matrix A has an eigenvalue 0, then A is singular.

Proof: (1) If $\lambda=0$ is an eigenvalue, then $K(\lambda) = \det(A-\lambda I)=0$ when $\lambda = 0$, which says $\det A = 0$, which says A is singular. (2) If $\lambda=0$ is an eigenvalue, then $A\mathbf{x} = \lambda\mathbf{x}$ for an eigenvector says $Ax = 0$, so the nullspace has dim > 0, so rank < N, so A is singular.

Definition: The set of n x n matrices A forms a vector space of dimension n^2, and the basis "vectors" can be taken to be the n^2 matrices with a 1 in a solitary position. Let's call these basis matrices M_i where i runs from 1 to n^2 and we agree to some order for them. A vector space must indicate how to apply a scalar to a vector, and in this space we have that if $\mathbf{A} = \{A_{ik}\}$, then $\alpha\mathbf{A} = \{\alpha A_{ik}\}$, so we just multiply each matrix element by α. See Stak p 96 for the definition of a vector space. Let's call this vector space the **Square Matrix Vector Space**.

Notice that we are using to thinking of a matrix as $A:E^n \rightarrow E^n$, where a matrix acts on column vectors. Here however we are doing $A:E^{n*n} \rightarrow E^{n*n}$. We can still think of matrices as the operators, but the things acted upon are now matrices. Probably $E^{n*n} = E^n \text{ x } E^n$ in some sense, but let's now worry about that interpretation right now.

Definition: Suppose we have a polynomial in λ *with matrix coefficients*, like $Q(\lambda) = B_o + B_1\lambda + B_2\lambda^2 + B_3\lambda^3$. In each term, since λ is a scalar, we can write λ^k before or after the matrix (or a mixture), because a scalar commutes with any matrix. However, when we talk about **extending a polynomial to a matrix argument**, we avoid ambiguity by requiring that the powers λ^k be to the right of the matrix coefficients in the polynomial, and only then do we replace λ with A. Remember that A probably does not commute with those matrix coefficients. Note that both $Q(\lambda)$ and Q(A) are matrices and are thus both elements in the vector space described above.

Counter Theorem: Let $C(\lambda) = A(\lambda)B(\lambda)$ where $A(\lambda)$ and $B(\lambda)$ are polynomials with matrix coefficients. Then if we extend all three polynomials to matrix arguments as described above, it is NOT true that $C(D) = A(D)B(D)$.

Proof: Suppose $A(\lambda) = A_0 + A_1\lambda$ and $B(\lambda) = B_0 + B_1\lambda$. Then

$C(\lambda) = A_0B_0 + (A_1B_0 + A_0B_1)\lambda + A_1B_1\lambda^2$, and then

$C(D) = C(\lambda) = A_0B_0 + (A_1B_0 + A_0B_1)D + A_1B_1D^2$. But

$A(D)B(D) = (A_0 + A_1D)(B_0 + B_1D) = A_0B_0 + A_1DB_0 + A_0B_1D + A_1DB_1D$.

Subtracting we find $C(D) - A(D)B(D) = A_1 (B_0D - DB_0) + A_1 (B_1D - DB_1)D = A_1[B_0, D] + A_1[B_1,D]D$. Since these commutators don't normally vanish, we find that normally $C(D) \neq A(D)B(D)$.

Theorem M41: Let $C(\lambda) = A(\lambda)B(\lambda)$ where $A(\lambda)$ and $B(\lambda)$ are polynomials with matrix coefficients. Then if $B(D) = 0$, it follows that $C(D) = 0$ [and this is true even though we know that $C(D) \neq A(D)B(D)$]

Proof: Write $C(D) = \Sigma_{jl} A_j B_k D^{j+k}$ where we have done or extension properly, that is, we write out $C(\lambda)$ as a polynomial in λ and *then* we substitute in D. We can rewrite the sum as $\Sigma_j A_j (\Sigma_k B_k D^k) D^j$. Now if $B(D) = 0$, then the inner sum is zero, and that then makes the whole thing be 0. QED.

Counter Theorem: Let $C(\lambda) = A(\lambda)B(\lambda)$ where $A(\lambda)$ and $B(\lambda)$ are polynomials with matrix coefficients. Then if $A(D) = 0$, it *does not follow* that $C(D) = 0$.

Proof: As before, $C(D) = \Sigma_{jl} A_j B_k D^{j+k}$. But now we cannot "get" a factor of D^j next to A_j because B is in the way. Thus, the proof as done above does not work here.

Theorem M42: Let $X(\lambda) = A(\lambda)B(\lambda).....F(\lambda)$ where all are polys with matrix coefficients. Then

(a) It is NOT true that X(D) = A(D)B(D).....F(D) for any matrix D.
(b) It is true that F(D) = 0 => X(D) = 0.

Proof: Just group the first set of polys together into some G(λ) so we have X(λ) = G(λ)F(λ). Then apply the previous theorems.

Corollary M42.1: Suppose P(λ) = Q(λ)(A - λI) where P and Q are polynomials in λ but with square matrix coefficients, and of course (A - λI) is a matrix as is A. If we then construct P(A) by extending P(λ) to a matrix argument as described above, then we find that P(A) = 0.

Proof: This is just an application of our C(λ) =A(λ)B(λ) theorem above where B(λ) = A - Iλ . If B(A) = 0, then we know that C(A) = 0. And it is clear that B(A) = 0.

Theorem M43 (**Cayley-Hamilton**). If K(λ) = 0 is the secular equation for matrix A, then K(A) = 0. This says that every matrix A satisfies its own secular equation.

Proof: We know that $[\text{cof}(B)]^T A = \det(B)\, I$ for any square matrix B. So let B = A - Iλ and then we have $\det(A - I\lambda) * I = [\text{cof}(A - I\lambda)]^T (A - I\lambda)$.

Now the matrix *elements* of $[\text{cof}(A - I\lambda)]^T$ are polynomials in λ, certainly of degree less than N, probably degree N-1, but the degree does not matter. We can expand this matrix using the basis matrices M_i of our Square Matrix Vector Space defined above, and the coefficient of each matrix M_i will be a polynomial in λ. We can then rearrange to get a polynomial in λ with matrix coefficients of some sort. It will be a mess, but all we care about is that $[\text{cof}(A - I\lambda)]^T$ is a polynomial in λ with matrix coefficients.

Clearly (A - Iλ) us a poly in λ with matrix coefficients.

What about the left side of our equation? It is K(λ)*I. where K(λ) = det(A - Iλ) = a normal polynomial, ie, one with numbers as coefficients. We can just slide I into each term and then say that the quantity K(λ)*I is then a polynomial in λ with matrix coefficients, and we might use some different letter to show this, like k(λ) = K(λ)*I .

So the upshot is that our equation $\det(A - I\lambda) * I = [\text{cof}(A - I\lambda)]^T (A - I\lambda)$ is of the form k(λ) =X(λ)Y(λ) and we use the previous theorem to say that Y(A) = 0 => k(A) = 0. Remember that this is true even though k(A) ≠ X(A)Y(A). Obviously Y(A) = 0 since Y(λ) = A - Iλ. Thus k(A) = 0, and we have then proven that a matrix satisfies its own secular equation, QED.

Corollary to Cayley-Hamilton. For any NxN matrix A, any positive power of A can be written as a linear combination of the N matrices A^0= I, A^1 = A, A^2, A^3, ...A^{N-1}. If A is invertible, then the conclusion also applies to negative powers of A. [Note: there may be some r < N for which this is also true.]

Proof: We know that k(λ) is monic and of degree N. That is, $k(\lambda) = I\,\lambda^N$ + lower power terms. Cayley-Hamilton shown that k(A) = A^N + lower power terms = 0. Thus we can write A^N = lincom of lower powers. Any higher power can then be broken down into the lower powers by doing it one power at a time. For example,

$$A^{N+2} = A^2 A^N = A^2 (aA^{N-1} + ...) = A (aA^N + ...) = A(a(aA^{N-1} + ...) + ...) = A (a^2 A^{N-1} ...)$$
$$= (a^2 A^N ...) = (a^2 (aA^{N-1} + ...) ...) = a^3 A^{N-1} + ...$$

Better Proof: We are interested in A^p for some positive power p. Start off and divide λ^p by $K(\lambda)$ to get $\lambda^p = q(\lambda)K(\lambda) + r(\lambda)$ where $r(\lambda)$ is poly of degree p-N. If this is > N, write $r(\lambda) = q'(\lambda)K(\lambda) + r'(\lambda)$ where now $r'(\lambda)$ is degree p - 2N. Keep going until you cannot go any more. The net result is that we can write $\lambda^p = Q(\lambda)K(\lambda) + R(\lambda)$ where $R(\lambda)$ is degree < N. Now apply our earlier theorem to the product $Q(\lambda)K(\lambda)$ to conclude that if $K(A) = 0$, then this product is 0. We then get $A^p = R(A)$, which is a poly in A of degree < N, QED for positive powers.

But now suppose $p < 0$. The write $\lambda^{-p} = Q(\lambda)K(\lambda) + R(\lambda)$ in the same manner. As before, we can kill the first term when $\lambda = A$, and we get $A^{-p} = R(\lambda)$, so A^{-p} must then *also* be a linear combination of the A^k for k = 0 to N-1. The meaning of A^{-p} is $(A^{-1})^p$ where A^{-1} is the inverse. This must exist for this negative power extension of our theorem to apply.

****

Sylvester's Law of Inertia

{ From the web: "The **index** of a matrix is characterized as the order of the largest Jordan block with zero eigenvalues. " But I don't think that is how I used it in my notes on Sylvester inertia. }

The following is based on information in linearAlg.pdf which is a good set of notes I downloaded:
www-local.cse.buffalo.edu/.../SUBNET/seminars/ InterconnectionNet_Spring02/papers/GeneralMath/linearAlg.ps

Def: The **inertia** of a matrix is the number of positive, negative and zero eigenvalues. It is just a triplet of numbers, like (3,2,1). Let's call this (N_+, N-, N_0), and $N_+ + N- + N_0 = N$, matrix size.

Def: The **signature** is (N_+ - N-) = s

Note that (N_+ + N-) = N - N_0 = rank = r.

{ Def: In my own notes, **index** = N_+]

Theorem M44: $N_0 = N - r$; $N_+ = (r+s)/2$, $N_- = (r-s)/2$.
If two matrices are related by a non-singular congruence, they have the same N and r.
Thus, (same signature) $\Leftrightarrow$ (same inertia).

Note: all these forms of Sylvester's law of inertia assume that matrices invovled are symmetric (or Hermitian), so that we know the eigenvalues are real.

Theorem M45: (**Sylvester's Law of Inertia**) Signature is invariant under a congruence transformation.
My PDF notes state it as : (A is congruent with B) $\Leftrightarrow$ (A and B have the same inertia).
The inertia of a symmetric matrix is not affected by a congruence transformation on that matrix.

Comments: It is assumed up front that A and B are symmetric so they have real eigenvalues, otherwise none of this makes sense. We know we can then diagonalize A to Λ_A. We could add a permutation to get the eigenvalues in the order positive, negative, then zero. We could then apply a diagonal congruence with all positive elements to get the matrix to be in **canonical form** which is diag (1,1,1... -1,-1,...,0,0..)

The total matrix to do this is then $S = CPR$ for Congruence, Permutation, Rotation. All three component matrices are non-singular, so S is invertible.

Suppose A and B have the same "inertia". Then you have $S_A^T A S_A = \text{Inertia} = S_B^T B S_B$ with the two S's as given above. Therefore, we can say that $A = S^T B S$ where $S = S_B (S_A)^{-1}$. This shows that two symmetric matrices with the same inertia must be congruent.

The other direction is harder. The idea is to let $a = N_+(A) =$ an integer, and think of the basis vectors $e_1 \ldots e_a$ which correspond to the positive eigenvalues (think of these as the eigenvectors of those eigenvalues). If you consider any x that is a lincom of these basis vectors, you will find that $(x,Ax) > 0$. I think the set of all such x form a subspace with dim = a of the overall space V. Now if A and B are congruent, then $A = S^T B S$, and we then have $(y,By) > 0$ for all y of the form $y = Sx$. This transform might kill some dimensions (ie, might have a non-trivial nullspace), so we conclude that $b \leq a$ where b is the dimension of the subspace of y where $(y,By) > 0$. But S is invertible so you can make the reverse argument to conclude that $a \leq b$, and therefore $a = b$. This means that for the size of the space spanned by the eigenvectors with positive eigenvalues for either A or B is the same. Since rank is the same, we know that N_0 are the same, thus the entire inertia must be the same. My argument may lack precision, but that is the gist of the argument.