MATH 2230 FINAL, FALL 2014

1. Let U be the set of $n \times n$ invertible real matrices (an open set inside the vector space of all $n \times n$ matrices), and B be a diagonal matrix with $\det B \neq 0$. Let $f : U \to U$ take $M \mapsto M^T BM$.

a [5]. Compute the derivative $Df|_M$.

Answer.

$$Df|_M(H) := \lim_{\varepsilon \to 0} \frac{(M + \varepsilon H)^T B(M + \varepsilon H)B - M^T BM}{\varepsilon} = \ldots = M^T BH + H^T BM$$

1b [10]. Show $f^{-1}(B)$ is a manifold.

Answer. $f(M)^T = f(M)$, so the derivative of f isn't going to be onto. Instead, define $g : U \to \{\text{symmetric matrices}\}$ by the same formula, so $f^{-1}(B) = g^{-1}(B)$. Now we need to check that the derivative of g is onto (i.e. onto the symmetric matrices) for all $M \in g^{-1}(B)$.

Say S is symmetric. If we can get $M^T BH = S/2$, then by transposing we get $H^T B^T M = S^T/2$ or $H^T BM = S/2$, so their sum will be S. So take $H = B^{-1}(M^T)^{-1}S/2$, and we can make these inverses since $\det B \neq 0$ and $M \in U$.

Now Dg is onto for all $M \in g^{-1}(B)$, so $g^{-1}(B)$ is a manifold.

(It's called "the orthogonal group of the form B", by the way. For example if $B = \mathrm{diag}(1, 1, 1, -1)$ then $g^{-1}(B)$ is the "Lorentz group" of rotations/reflections of spacetime.)

---

2. Let $M = \begin{pmatrix} 3 & a \\ 0 & b \end{pmatrix}$ where $a, b \in \mathbb{R}$.

a [5]. What are the eigenvalues of M?

Answer. Since this is upper triangular, the eigenvalues are on the diagonal: 3 and b.

2b [10]. Recall an **eigenbasis** for an $n \times n$ matrix M is an ordered list of vectors, each an eigenvector, forming a basis of $\mathbb{R}^n$ (or possibly $\mathbb{C}^n$, depending on context).

For which choices of a, b does $M = \begin{pmatrix} 3 & a \\ 0 & b \end{pmatrix}$ have an eigenbasis? Prove "for so-and-so pairs (a, b), M does have an eigenbasis; for all other pairs it does not."

Answer. If $b \neq 3$, then this has distinct eigenvalues, so it has an eigenbasis.

If $b = 3$ and $a = 0$, then this is $3I$, and the standard basis (or any other) is an eigenbasis.

If $b = 3$ and $a \neq 0$, then we try to solve $M\vec{v} = 3\vec{v}$, and discover that the only eigenvectors are multiples of $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$, i.e. we can't get two linearly independent eigenvectors.

---

3 [5]. You overhear Alice asking Bob about a matrix R, that she knows to be $k \times n$.

Bob: "R and R^T are both in echelon form."
Alice: "So what! There are still a million things R could be!"
Bob: "Indeed you're right: *exactly* a million."

What does this tell you about k and n (and why)?

Answer. If R is the zero matrix, we know what it looks like.

If R^T is in echelon form and is not zero, then the first column is not zero. If R is in echelon form and the first column is not zero, then the first column is a 1 atop 0s. But since R^T is in echelon form too, the first row is a 1 left of 0s. If we rip that row and column off, the resulting matrix is again in echelon form both ways. Continue until we get to a zero matrix.

So Alice knows that R has an $r \times r$ identity matrix in the NW corner, and is 0 everywhere else. The biggest matrix that would fit is of size $\min(k, n)$. So $0 \leq r \leq \min(k, n)$. For there to be 10^6 possibilities, $\min(k, n) = 10^6 - 1$.

4a [5]. Let $M = [\lambda]$, the 1×1 matrix. How many real, **orthonormal** eigenbases does M have? Your answer might depend on λ.

Answer. Psych! It doesn't depend on λ. Any vector $[m] \neq \vec{0}$ is an eigenvector, so the list $([m])$ is a 1-vector eigenbasis. To be orthonormal, though, we need $|[m]| = 1$, giving the $\text{\Huge two}$ possibilities $m = \pm 1$.

[In particular, if $\lambda = 0$, this is still true. It's fine for the *eigenvalue* to be 0, just not the *eigenvector*.

If you knew the theorem that *any* real symmetric matrix has an orthonormal eigenbasis, then you shouldn't have said "...except this one, with its 0 eigenvalue."]

4b [5]. Same question, for $M = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$.

Answer. Once again, every nonzero vector is an eigenvector (with eigenvalue -1). So any orthonormal basis $\left(\begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}, \begin{bmatrix} \mp\sin\theta \\ \pm\cos\theta \end{bmatrix}\right)$ will do, of which there are $\text{\Huge infinitely many.}$

4c [10]. Same question, for M a real symmetric matrix with eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$. (Do remember that an eigenbasis is an *ordered* list of vectors.)
Your answer might depend on $\lambda_1, \ldots, \lambda_n$.

Answer. If all the λ_i are distinct (as in (a)), then each eigenspace is 1-dimensional, and in that line there are two unit vectors. So choose one in each, and put them in order, for $2^n n!$ possibilities.

If some λ_i repeat (as in (b)), then that eigenspace has dimension > 1, and we can rotate our eigenvectors as in (b), for infinitely many possibilities.

---

5. Let M be the graph $\{(x, y, z) : z = xy\} \subset \mathbb{R}^3$, and N the graph $\{(x, y, z) : z = 0\} \subset \mathbb{R}^3$.

a [5]. For each point $p \in M \cap N$, compute the tangent spaces $T_p M$, $T_p N$.

Answer. You can either see M as the image of $p : \mathbb{R}^2 \to \mathbb{R}^3$, and compute the image of the derivative $Dp = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ y & x \end{pmatrix}$, or see it as $g^{-1}(0)$ where $g(x, y, z) = z - xy$, and compute the kernel of $Dg = \begin{pmatrix} -y & -x & 1 \end{pmatrix}$.

N you should just say "is a linear subspace, so all the tangent spaces are just N, again". I.e. the kernel of $\begin{pmatrix} 0 & 0 & 1 \end{pmatrix}$.

5b [5]. For which such p is $T_pM \neq T_pN$? What does this tell you about $M \cap N$, and why?

Answer. They are unequal for p on an axis, but not at the origin.

Where they are unequal, the derivative of the function $\mathbb{R}^3 \to \mathbb{R}^2$, $(x, y, z) \mapsto (z - xy, z)$ has rank 2, so is onto. Hence the intersection is a manifold there, *away* from the origin.

[Many people guessed something amounting to the reverse, that only at the origin is $M \cap N$ a manifold. That's the *bad* point.]

---

6 [15]. Let M be an invertible $n \times n$ matrix.
Show that for any open set $U \subseteq \mathbb{R}^n$, the image $M(U) := \{M\vec{u} : \vec{u} \in U\}$ is also an open set.
Hint: if this property holds for two matrices A, B, then it obviously holds for AB.

Answer. [One person gave a really cheap answer to this: "use the inverse function theorem!" Grudging respect.]

If U is itself a ball $B_r(\vec{u}) = \vec{u} + B_r(\vec{0})$, then $MU = M\vec{u} + MB_r(\vec{0})$. So we want to know, in particular, that $MB_r(\vec{0})$ contains a ball around $\vec{0}$.

Conversely, say we knew that $MB_r(\vec{0})$ contains a ball $B_{r'}(\vec{0})$. Then to show MU is open, we start with $\vec{v} \in MU$, find a ball of radius s around $M^{-1}\vec{v} \in U$, and learn that $s/rB_{r'}(\vec{v}) \subseteq MU$. So it's enough to know that $MB_r(\vec{0})$ contains a ball $B_{r'}(\vec{0})$.

Or, it might be easier to show that every $NB_r(\vec{0})$ is contained in a ball $B_{r''}(\vec{0})$, and then take $M = N^{-1}$. [This isn't hard to do directly using the matrix norm of N, which we didn't talk about.]

Every invertible M is a product of elementary matrices. So it's enough to check for M an elementary matrix.

If M switches two rows, then $MB_r(\vec{0}) = B_r(\vec{0})$, so that's easy.

If M scales a row by $\lambda \neq 0$, then $MB_r(\vec{0}) \supseteq B_{\min(r, r|\lambda|)}(\vec{0})$.

If M adds t times row i to row j, for $(v_1, \ldots, v_n) \in B_1(\vec{0})$, we get $|M\vec{v}|^2 = \sum_{k \neq j} v_k^2 + (v_i + tv_j)^2 \leq (n-1) + (1+|t|)^2$. So $MB_1(\vec{0}) \subseteq B_{\sqrt{n-1+(1+|t|)^2}}(\vec{0})$.

---

7. Let M(t) be a $k \times n$ matrix, where each entry is a polynomial in t, not just a number. Assume M(t) is not the zero matrix for any value of t.

Let $c(t) \in \{1, \ldots, n\}$ be the position of the leftmost pivot column of M(t).

a [1]. Why did we assume M(t) is not the zero matrix?

Answer. Otherwise there wouldn't be a "left pivot column".

7b [9]. Prove that there is a finite set $\{t_1, \ldots, t_d\}$ of t such that, for t not in that set, c(t) is some constant C.

Answer. Let C be the leftmost column of M that isn't all 0. Pick a nonzero entry p(t) in that column. Then for $p(t) \neq 0$, the matrix M(t) has C as its leftmost nonzero column, so $c(t) = C$. Let $\{t_1, \ldots, t_d\}$ be the places where $p(t) = 0$.

[Many people said "a polynomial has only finitely many roots." The correct statement is "a *nonzero* polynomial has only finitely many roots."]

---

8 [10]. Let M be an $n \times n$ matrix, and pick bases of $\mathit{ker}\ M$, $\mathit{image}\ M$. Assume that the bases, concatenated together, span $\mathbb{R}^n$.

Show that for $\vec{v} \in \mathbb{R}^n$, $M^2\vec{v} = \vec{0} \implies M\vec{v} = \vec{0}$.

Answer. If $\mathit{ker}\ M$ has dimension k, then $\mathit{image}\ M$ has dimension $n-k$, by the nullity+rank theorem. So together they're a list of length n. If that list spans, then it's linearly independent.

If $M^2\vec{v} = \vec{0}$, then $M\vec{v} \in \mathit{ker}\ M$. But $M\vec{v} \in \mathit{image}\ M$, by definition of image. So $M\vec{v} \in \mathit{ker}\ M \cap \mathit{image}\ M$, which is the zero subspace by the linear independence above.

[Examples: M the 0 matrix, M the identity matrix.]