MA 1B RECITATION 01/14/15

1. Warm-up Question

Consider the vector space of 3-by-3 matrices and consider the matrices in it that are "magic squares," where the rows, columns, and diagonals all add up to the same number, for example

$$\begin{bmatrix} 6 & 1 & 8 \\ 7 & 5 & 3 \\ 2 & 9 & 4 \end{bmatrix}.$$

We can get other easy examples of magic squares by rotating the above matrix, or by considering things like

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Do magic squares form a subspace? It turns out they do! You can easily check that the sum of magic squares is another magic square, and that scaling a magic square by some constant gives you another magic square. Another easy result is that the magic squares where all the rows, columns, diagonals sum to m form a subspace if and only if $m = 0$.

The dimension of 1-by-1 matrices is obviously 1, and the dimension of 2-by-2 matrices is 1 as well, with a basis given by

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

Now, what is the dimension of the space of 3×3 magic squares?

Bonus: What is the dimension of the space of $n \times n$ magic squares for $n > 3$?

2. Quick Recap

The first couple weeks of learning linear algebra just consists of definition after definition. Since definitions are the core of the course, it doesn't hurt to see them until you have them down cold. Let's quickly recall the basic facts and definitions about dimension and vector spaces.

We say a vector w is a **linear combination** of vectors $v_1, \ldots, v_n$ if there exists scalars $a_1, \ldots, a_n$ such that

$$w = a_1 v_1 + \cdots + a_n v_n.$$

The **span** of vectors $v_1, \ldots, v_n$, denoted $\operatorname{span}(v_1, \ldots, v_n)$ is the set of all vectors which are linear combinations of $v_1, \ldots, v_n$, that is,

$$\operatorname{span}(v_1, \ldots, v_n) = L(\{v_1, \ldots, v_n\}) = \{a_1 v_1 + \cdots a_n v_n \mid a_i \in \mathbf{R}\}.$$

Equivalently, the span is the smallest subspace containing $v_1, \ldots, v_n$.

Date: January 14, 2015.

Given a vector space V, if $\text{span}(v_1, \dots, v_n) = V$, we say that the vectors $v_1, \dots, v_n$ **span** V or **generate** V or that $\{v_1, \dots, v_n\}$ is a **spanning** or **generating set** for V.

We say that the vectors $v_1, \dots, v_n$ are **linearly independent** if

$$a_1 v_1 + \cdots a_n v_n = 0$$

implies that $a_1 = \cdots = a_n = 0$. (Otherwise, we say that $v_1, \dots, v_n$ are linearly *dependent*.)

Proposition 1. *The vectors $v_1, \dots, v_n$ are linearly independent if and only if every vector $w \in span(v_1, \dots, v_n)$ can be written **uniquely** as*

$$w = a_1 v_1 + \cdots + a_n v_n.$$

A linearly independent spanning set of vectors for V is a **basis** for V.

Here are some basic facts about bases, some of which we went over in lecture, some of which you will prove on the homework.

- Every vector space has a basis.
- Any generating set for V contains a basis.
- Any linearly independent set of vectors can be extended to a basis.
- Any two bases of V contain the same number of elements.

3. Dimension

Since any two bases of V contain the same number of elements, we say that a vector space V is of **dimension** n if its basis consists of n elements.

Here are a couple of obvious consequences of the definition.

- The dimension of $\mathbf{R}^n$ is n.
- The dimension of a vector space is a nonnegative integer, and there exists a vector space of dimension n for every nonnegative integer n.
- The dimension of a *proper* subspace of a finite-dimensional vector space is less than the dimension of the whole space.
- If A is a set of m vectors in V and $m < \dim(V)$, then A does not contain a basis.
- If A is a set of m vectors in V, and $m > \dim(V)$, then A is linearly dependent.

Dimension in linear algebra is a pretty "coarse" invariant, but it is pretty much the only invariant of vector spaces that we care about! This is because *any* real (or complex) vector space of dimension n is isomorphic (as a vector space) to $\mathbf{R}^n$, that is, they are essentially the same in the eyes of linear algebra. This is a no-brainer for something like $\mathbf{R}^n$, but, of course, vector spaces can look very different upon first glance!

Here are a couple of interesting mathematical objects that are "the same" as vector spaces. (The precise term for this is that the vector spaces are *isomorphic*, but we will touch on that slightly later in the course.)

Example 1. Consider the set $\mathbf{R}^4$ and $M_2(\mathbf{R})$ the vector space of 2-by-2 real matrices. They are both 4-dimensional over $\mathbf{R}$

Example 2. Consider the vector spaces of $\mathbf{R}^2$ and $\mathbf{C}$, considered as a real vector space (that is, forget that $\mathbf{C}$ has a magical multiplication). Then both $\mathbf{R}^2$ and $\mathbf{C}$ have dimension 2 over $\mathbf{R}$ and so are isomorphic as vector spaces.

Example 3. If V is a real vector space of dimension n, then $\mathcal{L}(V, \mathbf{R})$, the set of linear maps from V to $\mathbf{R}$ is also of dimension n, so the vector space and the real-valued linear maps are "the same" as vector spaces.

4. Answer to the magic squares question

This is a heuristic solution and is missing many details, but you can turn this reasoning into a rigorous proof with some work.

In general a magic square can be considered as an n^2-dimensional vector that satisfies $n+n+2-1-1$ independent constraints: the rows, columns, and diagonals have the same sum, but the sum of the row equations equals the sum of the column equations, so one of those is redundant.

Thus, it turns out the magic squares where each sum to zero have dimension $n^2 - 2n - 1$ for $n \geq 3$.

To a basis for the full set of magic squares (where the sum can be anything), we can take the basis above and adding the element that consists of 1's in all places.

Thus, if $n \geq 3$, the dimension is $n^2 - 2n$.

For instance, every 3-by-3 magic square can be written as

$$\begin{bmatrix} a+b & a+c-b & a-c \\ a-b-c & a & a+b+c \\ a+c & a+b-c & a-b \end{bmatrix}.$$

5. Big Picture: A Classification of Linear Transformations, Affine Transformations

What does a linear transformation "look like"?

It turns out that there is a one-to-one correspondence between linear transformations $T : \mathbf{R}^n \to \mathbf{R}^m$ and $m \times n$ matrices. This is one of the major conceptual milestones of this course.

For example, for linear maps $\mathbf{R}^2 \to \mathbf{R}^2$, all of them are of the form

$$\begin{aligned} T : \mathbf{R}^2 &\to \mathbf{R}^2 \\ (x, y) &\mapsto (ax + by, cx + dy) \end{aligned}$$

which corresponds to the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$. Thus, these linear maps have a similar shape, namely, they correspond to a scaling or "shearing" operation that you can visualize by say, fixing a unit square with lower-left corner at $(0, 0)$ and seeing what happens to the square when you let T act on $\mathbf{R}^2$.

However, this is not the complete truth. I've implicitly made an assumption about choices of bases. Namely, I've assumed that both copies of $\mathbf{R}^2$ have the same basepoint (the origin 0) and the same bases x and y, corresponding to the standard basis $e_1 = (1, 0)$ and $e_2 = (0, 1)$.

We could also suppose that the second copy of $\mathbf{R}^2$ has a different basepoint, say, $(3, 1)$ instead of the $(0, 0)$, and choose bases x' and y' *relative to the basepoint*, so x' is given by the vector $(3, 1) \to (4, 1)$ and y' is given by the vector $(3, 1) \to (3, 2)$. To distinguish this slightly modified copy, let's call this $\widetilde{\mathbf{R}^2}$. This is still a two-dimensional vector space and so is isomorphic to $\mathbf{R}^2$. Since linear maps must map the origin to origin, the linear maps $\mathbf{R}^2 \to \mathbf{R}^2$ are of the form

$$\begin{aligned} T : \mathbf{R}^2 &\to \widetilde{\mathbf{R}^2} \\ (x, y) &\mapsto (ax' + by' + 3, cx' + dy' + 1) \end{aligned}$$

We see that now we have *translation* by $(3, 1)$ as a valid map from $\mathbf{R}^2$ to a copy of $\mathbf{R}^2$. Thus, if we fix the same basepoint and bases, a translation is *not* a linear transformation $\mathbf{R}^2 \to \mathbf{R}^2$, but if we allow ourselves to change the basepoint, we allow for translations.

Now, when we usually talk about maps $\mathbf{R}^2 \to \mathbf{R}^2$, we'll assume they have the same basepoints for simplicity, but maybe you're wondering, if we give ourselves the freedom to change basepoints, does that preserve anything about, say, lines or shapes in the plane? Or is it completely wild?

It turns out that it does in fact preserve a lot of information. If we allow ourselves to change basepoints, we get a class of maps called *affine transformations* that give us all the ways we can "transform" $\mathbf{R}^2$ that preserves straight lines and ratios of distances between points lying on a straight line. Such maps don't necessarily preserve angles or lengths, but many things do hold, like points lying on lines will still remain on the line after transformation, midpoints of line segments will remain midpoints of that line segment after transformation, and parallel lines will remain parallel after transformation.

In the language of linear algebra, we can describe affine transformations very concisely: they are merely linear transformations followed by a translation. Do affine transformations form a vector space? It turns out they do! What is the dimension of said vector space? Note that the linear transformations are a proper subspace of the vector space of affine transformations.

But why do we care about these weird invariant properties? It turns out that they can have a lot of scientific meaning. For instance, results like Noether's theorem in physics says that invariants are a sign of symmetry, and can point to the fact that seemingly complicated phenomena in the world can admit a surprisingly simple description if you look at the problem the right way.

We will stick to studying linear transformation in this Ma 1b, but you will surely encounter affine transformations in later courses.

6. Analyzing Linear Transformations

Recall that we have the rank-nullity theorem, one of the few formulas in this class: given a linear transformation $T : V \to W$, we have

$$\dim V = \dim \operatorname{Ker}(T) + \dim \operatorname{Im}(T) = \text{"nullity"} + \text{"rank"},$$

where $\operatorname{Ker}(T) = \{v \in V \mid Tv = 0\}$ and $\operatorname{Im}(T) = \{w \in W \mid Tv = w \text{ for some } v \in V\}$. In matrix form, this means that rank plus nullity equals the number of columns.

Example 1. Let's consider the "crushing" map $T : \mathbf{R}^2 \to \mathbf{R}^2$ given by

$$(x, y) \mapsto (x, 0).$$

To begin, let's explore how this transforms some shapes in $\mathbf{R}^2$. For instance, it turns a unit square into a line. It turns a mountain shape into a line.

What does it look like in matrix form? Let's consider it with the standard bases $e_1 = (1,0)^T$ and $e_2 = (0,1)^T$ on each side. It is

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

and we can check, since (x, y) corresponds to $xe_1 + ye_2$, we have

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x \\ 0 \end{bmatrix}.$$

Great!

What's the rank and nullity? The elements of $\mathbf{R}^2$ that map to 0 under T are all of the form $(0, a)^T$ where $a \in \mathbf{R}$, so it is spanned, for instance, by $(0, 1)^T$.

Note that the image of T consists of elements of the form $(a, 0)$ with $a \in \mathbf{R}$ so is spanned, for instance, by $(1, 0)$. Thus, we see that the rank-nullity theorem holds here:

$$2 = 1 + 1.$$

As you can see, the rank and nullity are somewhat intuitive. Roughly speaking, the nullity represents the dimensions that are "lost" in the transformation, and the rank represents those that are "preserved," and the sum of the dimensions that we "lose" and "preserve" must be just what we started with.