Several variables

Often we have not only one, but several variables in a problem. The issues that come up are somewhat more complex than for one variable. Let us first start with vector spaces and linear functions.

While it is common to use $\vec{x}$ or $\mathbf{x}$ for elements of $\mathbb{R}^n$, especially in the applied sciences, we will just use x, which is common in mathematics. That is $x \in \mathbb{R}^n$ is a *vector* which means that $x = (x^1, x^2, \dots, x^n)$ is an n-tuple of real numbers. We will use upper indices for identifying component. That leaves us the lower index for sequences of vectors. That is we can have vectors x_1 and x_2 in $\mathbb{R}^n$ and then $x_1 = (x_1^1, x_1^2, \dots, x_1^n)$ and $x_2 = (x_2^1, x_2^2, \dots, x_2^n)$. It is common to write vectors as *column vectors*, that is

$$x = (x^1, x^2, \dots, x^n) = \begin{bmatrix} x^1 \\ x^2 \\ \vdots \\ x^n \end{bmatrix},$$

and we will do so when convenient. We will use this notation with square brackets and use round brackets for just an n-tuple of numbers. We will call real numbers *scalars* to distinguish them from vectors.

Definition: A set X together with operations $+\colon X \times X \to X$ and $\cdot\colon \mathbb{R} \times X \to X$ (where we will just write ax instead of $a \cdot x$) is called a *vector space* (or a *real vector space*) if the following conditions are satisfied:

(i) (Addition is associative) If $u, v, w \in X$, then $u + (v + w) = (u + v) + w$.
(ii) (Addition is commutative) If $u, v \in X$, then $u + v = v + u$.
(iii) (Additive identity) There is a $0 \in X$ such that $v + 0 = v$ for all $v \in X$.
(iv) (Additive inverse) For every $v \in X$, there is a $-v \in X$, such that $v + (-v) = 0$.
(v) (Distributive law) If $a \in \mathbb{R}$, $u, v \in X$, then $a(u + v) = au + av$.
(vi) (Distributive law) If $a, b \in \mathbb{R}$, $v \in X$, then $(a + b)v = av + bv$.
(vii) (Multiplication is associative) If $a, b \in \mathbb{R}$, $v \in X$, then $(ab)v = a(bv)$.
(viii) (Multiplicative identity) $1v = v$ for all $v \in X$.

Elements of a vector space are usually called *vectors*, even if they are not elements of $\mathbb{R}^n$ (vectors in the "traditional" sense).

In short X is an Abelian group with respect to the addition, and equipped with multiplication by scalars.

An example vector space is $\mathbb{R}^n$, where addition and multiplication by a constant is done componentwise. We will mostly deal with vector spaces as subsets of $\mathbb{R}^n$, but there are other vector spaces that are useful in analysis. For example, the space $C([0,1], \mathbb{R})$ of continuous functions on the interval $[0,1]$ is a vector space. The functions $L^1(X, \mu)$ is also a vector space.

A trivial example of a vector space (the smallest one in fact) is just $X = \{0\}$. The operations are defined in the obvious way. You always need a zero vector to exist, so all vector spaces are nonempty sets.

It is also possible to use other fields than $\mathbb{R}$ in the definition (for example it is common to use $\mathbb{C}$), but let us stick with $\mathbb{R}$. [1]

Definition: If we have vectors $x_1, \dots, x_k \in \mathbb{R}^n$ and scalars $a^1, \dots, a^k \in \mathbb{R}$, then

$$a^1 x_1 + a^2 x_2 + \cdots + a^k x_k$$

is called a *linear combination* of the vectors $x_1, \dots, x_k$.

Note that if $x_1, \dots, x_k$ are in a vector space X, then any linear combination of $x_1, \dots, x_k$ is also in X.

If $Y \subset \mathbb{R}^n$ is a set then the *span* of Y, or in notation $\operatorname{span}(Y)$, is the set of all linear combinations of some finite number of elements of Y. We also say that Y *spans* $\operatorname{span}(Y)$.

Example: Let $Y = \{(1,1)\} \subset \mathbb{R}^2$. Then

$$\operatorname{span}(Y) = \{(x, x) \in \mathbb{R}^2 : x \in \mathbb{R}\},$$

or in other words the line through the origin and the point $(1,1)$.

[1] If you want a funky vector space over a different field, $\mathbb{R}$ is an infinite dimensional vector space over the rational numbers.

Example: Let $Y = \{(1,1),(0,1)\} \subset \mathbb{R}^2$. Then

$$\operatorname{span}(Y) = \mathbb{R}^2,$$

as any point $(x,y) \in \mathbb{R}^2$ can be written as a linear combination

$$(x,y) = x(1,1) + (y-x)(0,1).$$

Since a sum of two linear combinations is again a linear combination, and a scalar multiple of a linear combination is a linear combination, we see that:

Proposition: A span of a set $Y \subset \mathbb{R}^n$ is a *vector space*.

If Y is already a vector space then $\operatorname{span}(Y) = Y$.

Definition: A set of vectors $\{x_1, x_2, \ldots, x_k\}$ is said to be *linearly independent*, if the only solution to

$$a^1 x_1 + \cdots + a^k x_k = 0$$

is the trivial solution $a^1 = \cdots = a^k = 0$. Here 0 is the vector of all zeros. A set that is not linearly independent, is *linearly dependent.*

Note that if a set is linearly dependent, then this means that one of the vectors is a linear combination of the others.

A linearly independent set B of vectors that span a vector space X is called a *basis* of X.

If a vector space X contains a linearly independent set of d vectors, but no linearly independent set of $d+1$ vectors then we say the *dimension* or $\dim X = d$. If for all $d \in \mathbb{N}$ the vector space X contains a set of d linearly independent vectors, we say X is infinite dimensional and write $\dim X = \infty$.

Note that we have $\dim\{0\} = 0$. So far we have not shown that any other vector space has a finite dimension. We will see in a moment that any vector space that is a subset of $\mathbb{R}^n$ has a finite dimension, and that dimension is less than or equal to n.

Proposition: If $B = \{x_1, \ldots, x_k\}$ is a basis of a vector space X, then every point $y \in X$ has a unique representation of the form

$$y = \sum_{j=1}^{k} \alpha^j x_j$$

for some numbers $\alpha^1, \ldots, \alpha^k$.

Proof. First, every $y \in X$ is a linear combination of elements of B since X is the span of B. For uniqueness suppose

$$y = \sum_{j=1}^{k} \alpha^j x_j = \sum_{j=1}^{k} \beta^j x_j$$

then

$$\sum_{j=1}^{k} (\alpha^j - \beta^j) x_j = 0$$

so by linear independence of the basis $\alpha^j = \beta^j$ for all j. $\square$

For $\mathbb{R}^n$ we define

$$e_1 = (1,0,0,\ldots,0), \quad e_2 = (0,1,0,\ldots,0), \quad \ldots, \quad e_n = (0,0,0,\ldots,1),$$

and call this the *standard basis* of $\mathbb{R}^n$. We use the same letters e_j for any $\mathbb{R}^n$, and what space $\mathbb{R}^n$ we are working in is understood from context. A direct computation shows that $\{e_1, e_2, \ldots, e_n\}$ is really a basis of $\mathbb{R}^n$; it is easy to show that it spans $\mathbb{R}^n$ and is linearly independent. In fact,

$$x = (x^1, \ldots, x^n) = \sum_{j=1}^{n} x^j e_j.$$

Proposition (Theorems 9.2 and 9.3 in Rudin): Suppose that X is a vector space.

(i) If X is spanned by d vectors, then $\dim X \leq d$.
(ii) $\dim X = d$ if and only if X has a basis of d vectors (and so every basis has d vectors).
(iii) In particular, $\dim \mathbb{R}^n = n$.
(iv) If $Y \subset X$ is a vector space and $\dim X = d$, then $\dim Y \leq d$.
(v) If $\dim X = d$ and a set T of d vectors spans X, then T is linearly independent.
(vi) If $\dim X = d$ and a set T of m vectors is linearly independent, then there is a set S of $d - m$ vectors such that $T \cup S$ is a basis of X.

Proof. Let us start with (i). Suppose that $S = \{x_1, \dots, x_d\}$ span X. Now suppose that $T = \{y_1, \dots, y_m\}$ is a set of linearly independent vectors of X. We wish to show that $m \leq d$. Write

$$y_1 = \sum_{k=1}^{d} \alpha_1^k x_k,$$

which we can do as S spans X. One of the α_1^k is nonzero (otherwise y_1 would be zero), so suppose without loss of generality that this is α_1^1. Then we can solve

$$x_1 = \frac{1}{\alpha_1^1} y_1 - \sum_{k=2}^{d} \frac{\alpha_1^k}{\alpha_1^1} x_k.$$

In particular $\{y_1, x_2, \dots, x_d\}$ span X, since x_1 can be obtained from $\{y_1, x_2, \dots, x_d\}$. Next,

$$y_2 = \alpha_2^1 y_1 + \sum_{k=2}^{d} \alpha_2^k x_k,$$

As T is linearly independent, we must have that one of the α_2^k for $k \geq 2$ must be nonzero. Without loss of generality suppose that this is α_2^2. Proceed to solve for

$$x_2 = \frac{1}{\alpha_2^2} y_2 - \frac{\alpha_2^1}{\alpha_2^2} y_1 - \sum_{k=3}^{d} \frac{\alpha_2^k}{\alpha_2^2} x_k.$$

In particular $\{y_1, y_2, x_3, \dots, x_d\}$ spans X. The astute reader will think back to linear algebra and notice that we are row-reducing a matrix.

We continue this procedure. Either $m < d$ and we are done. So suppose that $m \geq d$. After d steps we obtain that $\{y_1, y_2, \dots, y_d\}$ spans X. So any other vector v in X is a linear combination of $\{y_1, y_2, \dots, y_d\}$, and hence cannot be in T as T is linearly independent. So $m = d$.

Let us look at (ii). First notice that if we have a set T of k linearly independent vectors that do not span X, then we can always choose a vector $v \in X \setminus \operatorname{span}(T)$. The set $T \cup \{v\}$ is linearly independent (exercise). If $\dim X = d$, then there must exist some linearly independent set of d vectors T, and it must span X, otherwise we could choose a larger set of linearly independent vectors. So we have a basis of d vectors. On the other hand if we have a basis of d vectors, it is linearly independent and spans X. By (i) we know there is no set of $d + 1$ linearly independent vectors, so dimension must be d.

For (iii) notice that $\{e_1, \dots, e_n\}$ is a basis of $\mathbb{R}^n$.

To see (iv), suppose that Y is a vector space and $Y \subset X$, where $\dim X = d$. As X cannot contain $d + 1$ linearly independent vectors, neither can Y.

For (v) suppose that T is a set of m vectors that is linearly dependent and spans X. Then one of the vectors is a linear combination of the others. Therefore if we remove it from T we obtain a set of $m - 1$ vectors that still span X and hence $\dim X \leq m - 1$.

For (vi) suppose that $T = \{x_1, \dots, x_m\}$ is a linearly independent set. We follow the procedure above in the proof of (ii) to keep adding vectors while keeping the set linearly independent. As the dimension is d we can add a vector exactly $d - m$ times. $\square$

Definition: A mapping $A\colon X \to Y$ of vector spaces X and Y is said to be *linear* (or a *linear transformation*) if for every $a \in \mathbb{R}$ and $x, y \in X$ we have

$$A(ax) = aA(x) \qquad A(x+y) = A(x) + A(y).$$

We will usually just write Ax instead of $A(x)$ if A is linear.

If A is one-to-one an onto then we say A is *invertible* and we define A^{-1} as the inverse.

If $A\colon X \to X$ is linear then we say A is a *linear operator* on X.

We will write $L(X,Y)$ for the set of all linear transformations from X to Y, and just $L(X)$ for the set of linear operators on X. If $a, b \in \mathbb{R}$ and $A, B \in L(X,Y)$ then define the transformation $aA + bB$

$$(aA + bB)(x) = aAx + bBx.$$

It is not hard to see that $aA + bB$ is linear.

If $A \in L(Y,Z)$ and $B \in L(X,Y)$, then define the transformation AB as

$$ABx = A(Bx).$$

It is trivial to see that $AB \in L(X,Z)$.

Finally denote by $I \in L(X)$ the *identity*, that is the linear operator such that $Ix = x$ for all x.

Note that it is obvious that $A0 = 0$.

Proposition: If $A\colon X \to Y$ is invertible, then A^{-1} is linear.

Proof. Let $a \in \mathbb{R}$ and $y \in Y$. As A is onto, then there is an x such that $y = Ax$, and further as it is also one-to-one $A^{-1}(Az) = z$ for all $z \in X$. So

$$A^{-1}(ay) = A^{-1}(aAx) = A^{-1}(A(ax)) = ax = aA^{-1}(y).$$

Similarly let $y_1, y_2 \in Y$, and $x_1, x_2 \in X$ such that $Ax_1 = y_1$ and $Ax_2 = y_2$, then

$$A^{-1}(y_1 + y_2) = A^{-1}(Ax_1 + Ax_2) = A^{-1}(A(x_1 + x_2)) = x_1 + x_2 = A^{-1}(y_1) + A^{-1}(y_2).$$

□

Proposition: If $A\colon X \to Y$ is linear then it is completely determined by its values on a basis of X. Furthermore, if B is a basis, then any function $\tilde{A}\colon B \to Y$ extends to a linear function on X.

Proof. For infinite dimensional spaces, the proof is essentially the same, but a little trickier to write, so let's stick with finitely many dimensions. Let $\{x_1, \ldots, x_n\}$ be a basis and suppose that $A(x_j) = y_j$. Then every $x \in X$ has a unique representation

$$x = \sum_{j=1}^{n} b^j x_j$$

for some numbers $b^1, \ldots, b^n$. Then by linearity

$$Ax = A \sum_{j=1}^{n} b^j x_j = \sum_{j=1}^{n} b^j Ax_j = \sum_{j=1}^{n} b^j y_j.$$

The "furthermore" follows by defining the extension $Ax = \sum_{j=1}^{n} b^j y_j$, and noting that this is well defined by uniqueness of the representation of x. □

Theorem 9.5: If X is a finite dimensional vector space and $A\colon X \to X$ is linear, then A is one-to-one if and only if it is onto.

Proof. Let $\{x_1, \ldots, x_n\}$ be a basis for X. Suppose that A is one-to-one. Now suppose

$$\sum_{j=1}^{n} c^j A x_j = A \sum_{j=1}^{n} c^j x_j = 0.$$

As A is one-to-one, the only vector that is taken to 0 is 0 itself. Hence,

$$0 = \sum_{j=1}^{n} c^j x_j$$

and so $c^j = 0$ for all j. Therefore, $\{Ax_1, \ldots, Ax_n\}$ is linearly independent. By an above proposition and the fact that the dimension is n, we have that $\{Ax_1, \ldots, Ax_n\}$ span X. As any point $x \in X$ can be written as

$$x = \sum_{j=1}^{n} a^j A x_j = A \sum_{j=1}^{n} a^j x_j,$$

so A is onto.

Now suppose that A is onto. As A is determined by the action on the basis we see that every element of X has to be in the span of $\{Ax_1, \ldots, Ax_n\}$. Suppose that

$$A \sum_{j=1}^{n} c^j x_j = \sum_{j=1}^{n} c^j A x_j = 0.$$

By the same proposition as $\{Ax_1, \ldots, Ax_n\}$ span X, the set is independent, and hence $c^j = 0$ for all j. This means that A is one-to-one. If $Ax = Ay$, then $A(x - y) = 0$ and so $x = y$. □

Let us start measuring distance. If X is a vector space, then we say a real valued function $\|\cdot\|$ is a *norm* if:

(i) $\|x\| \geq 0$, with $\|x\| = 0$ if and only if $x = 0$.
(ii) $\|cx\| = |c| \|x\|$ for all $c \in \mathbb{R}$ and $x \in X$.
(iii) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in X$.

Let $x = (x^1, \ldots, x^n) \in \mathbb{R}^n$. Define

$$\|x\| = \sqrt{(x^1)^2 + (x^2)^2 + \cdots + (x^n)^2}$$

be the *Euclidean norm*. Then $d(x, y) = \|x - y\|$ is the standard distance function on $\mathbb{R}^n$ that we used when we talked about metric spaces. In fact we proved last semester that the Euclidean norm is a norm. On any vector space X, once we have a norm, we can define a distance $d(x, y) = \|x - y\|$ that makes X a metric space.

Let $A \in L(X, Y)$. Define

$$\|A\| = \sup\{\|Ax\| : x \in X \quad \text{with} \quad \|x\| = 1\}.$$

The number $\|A\|$ is called the *operator norm*. We will see below that indeed it is a norm (at least for finite dimensional spaces). By linearity we get

$$\|A\| = \sup\{\|Ax\| : x \in X \quad \text{with} \quad \|x\| = 1\} = \sup_{\substack{x \in X \\ x \neq 0}} \frac{\|Ax\|}{\|x\|}.$$

This implies that

$$\|Ax\| \leq \|A\| \|x\| .$$

It is not hard to see from the definition that $\|A\| = 0$ if and only if $A = 0$, that is, if A takes every vector to the zero vector.

For finite dimensional spaces $\|A\|$ is always finite as we will prove below. For infinite dimensional spaces this need not be true. For a simple example, take the vector space of continuously differentiable functions on $[0, 1]$ and as the norm use the uniform norm. Then for example the functions $\sin(nx)$ have norm 1, but the derivatives have norm n. So differentiation (which is a linear operator) has unbounded norm on this space. But let us stick to finite dimensional spaces now.

Proposition (Theorem 9.7 in Rudin):

(i) If $A \in L(\mathbb{R}^n, \mathbb{R}^m)$, then $\|A\| < \infty$ and A is uniformly continuous (Lipschitz with constant $\|A\|$).

(ii) If $A, B \in L(\mathbb{R}^n, \mathbb{R}^m)$ and $c \in \mathbb{R}$, then

$$\|A + B\| \leq \|A\| + \|B\|, \qquad \|cA\| = |c| \, \|A\| .$$

In particular $L(\mathbb{R}^n, \mathbb{R}^m)$ is a metric space with distance $\|A - B\|$.

(iii) If $A \in L(\mathbb{R}^n, \mathbb{R}^m)$ and $B \in L(\mathbb{R}^m, \mathbb{R}^k)$, then

$$\|BA\| \leq \|B\| \, \|A\| .$$

Proof. For (i), let $x \in \mathbb{R}^n$. We know that A is defined by its action on a basis. Write

$$x = \sum_{j=1}^{n} c^j e_j .$$

Then

$$\|Ax\| = \left\| \sum_{j=1}^{n} c^j A e_j \right\| \leq \sum_{j=1}^{n} \left| c^j \right| \|A e_j\| .$$

If $\|x\| = 1$, then it is easy to see that $|c^j| \leq 1$ for all j, so

$$\|Ax\| \leq \sum_{j=1}^{n} \left| c^j \right| \|A e_j\| \leq \sum_{j=1}^{n} \|A e_j\| .$$

The right hand side does not depend on x and so we are done, we have found a finite upper bound. Next,

$$\|A(x-y)\| \leq \|A\| \, \|x - y\|$$

as we mentioned above. So if $\|A\| < \infty$, then this says that A is Lipschitz with constant $\|A\|$.

For (ii), let us note that

$$\|(A+B)x\| = \|Ax + Bx\| \leq \|Ax\| + \|Bx\| \leq \|A\| \, \|x\| + \|B\| \, \|x\| = (\|A\| + \|B\|) \, \|x\| .$$

So $\|A + B\| \leq \|A\| + \|B\|$. Similarly

$$\|(cA)x\| = |c| \, \|Ax\| \leq (|c| \, \|A\|) \, \|x\| .$$

So $\|cA\| \leq |c| \, \|A\|$. Next note

$$|c| \, \|Ax\| = \|cAx\| \leq \|cA\| \, \|x\| .$$

So $|c| \, \|A\| \leq \|cA\|$.

That we have a metric space follows pretty easily, and is left to student.

For (iii) write

$$\|BAx\| \leq \|B\| \, \|Ax\| \leq \|B\| \, \|A\| \, \|x\| .$$

$\square$

As a norm defines a metric, we have defined a metric space topology on $L(\mathbb{R}^n, \mathbb{R}^m)$ so we can talk about open/closed sets, continuity, and convergence. Note that we have defined a norm only on $\mathbb{R}^n$ and not on an arbitrary finite dimensional vector space. However, after picking bases, we can define a norm on any vector space in the same way. So we really have a topology on any $L(X, Y)$, although the precise metric would depend on the basis picked.

Theorem 9.8: Let $U \subset L(\mathbb{R}^n)$ be the set of invertible linear operators.

(i) If $A \in U$ and $B \in L(\mathbb{R}^n)$, and

$$\|A - B\| < \frac{1}{\|A^{-1}\|}, \tag{1}$$

then B is invertible.

(ii) U is open and $A \mapsto A^{-1}$ is a continuous function on U.

The theorem says that U is an open set and $A \mapsto A^{-1}$ is continuous on U.

You should always think back to $\mathbb{R}^1$, where linear operators are just numbers a. The operator a is invertible ($a^{-1} = 1/a$) whenever $a \neq 0$. Of course $a \mapsto 1/a$ is continuous. When $n > 1$, then there are other noninvertible operators, and in general things are a bit more difficult.

Proof. Let us prove (i). First a straight forward computation

$$\|x\| = \left\|A^{-1}Ax\right\| \leq \left\|A^{-1}\right\| \|Ax\| \leq \left\|A^{-1}\right\| (\|(A-B)x\| + \|Bx\|) \leq \left\|A^{-1}\right\| \|A-B\| \|x\| + \left\|A^{-1}\right\| \|Bx\| .$$

Now assume that $x \neq 0$ and so $\|x\| \neq 0$. Using (1) we obtain

$$\|x\| < \|x\| + \left\|A^{-1}\right\| \|Bx\| ,$$

or in other words $\|Bx\| \neq 0$ for all nonzero x, and hence $Bx \neq 0$ for all nonzero x. This is enough to see that B is one-to-one (if $Bx = By$, then $B(x-y) = 0$, so $x = y$). As B is one-to-one operator from $\mathbb{R}^n$ to $\mathbb{R}^n$ it is onto and hence invertible.

Let us look at (ii). Let B be invertible and near A^{-1}, that is (1) is satisfied. In fact, suppose that $\|A-B\| \|A^{-1}\| < 1/2$. Then we have shown above (using $B^{-1}y$ instead of x)

$$\left\|B^{-1}y\right\| \leq \left\|A^{-1}\right\| \|A-B\| \left\|B^{-1}y\right\| + \left\|A^{-1}\right\| \|y\| \leq 1/2 \left\|B^{-1}y\right\| + \left\|A^{-1}\right\| \|y\| ,$$

or

$$\left\|B^{-1}y\right\| \leq 2 \left\|A^{-1}\right\| \|y\| .$$

So $\|B^{-1}\| \leq 2 \|A^{-1}\|$.

Now note that

$$A^{-1}(A-B)B^{-1} = A^{-1}(AB^{-1} - I) = B^{-1} - A^{-1},$$

and

$$\left\|B^{-1} - A^{-1}\right\| = \left\|A^{-1}(A-B)B^{-1}\right\| \leq \left\|A^{-1}\right\| \|A-B\| \left\|B^{-1}\right\| \leq 2 \left\|A^{-1}\right\|^2 \|A-B\| .$$

□

Finally let's get to matrices, which are a convenient way to represent finite-dimensional operators. If we have bases $\{x_1, \ldots, x_n\}$ and $\{y_1, \ldots, y_m\}$ for vector spaces X and Y, then we know that a linear operator is determined by its values on the basis. Given $A \in L(X,Y)$, define the numbers $\{a_i^j\}$ as follows

$$Ax_j = \sum_{i=1}^{m} a_j^i y_i,$$

and write them as a *matrix*

$$A = \begin{bmatrix} a_1^1 & a_2^1 & \cdots & a_n^1 \\ a_1^2 & a_2^2 & \cdots & a_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^m & a_2^m & \cdots & a_n^m \end{bmatrix} .$$

Note that the *columns* of the matrix are precisely the coefficients that represent Ax_j. We can represent x_j as a column vector of n numbers (an $n \times 1$ matrix) with 1 in the jth position and zero elsewhere, and then

$$Ax = \begin{bmatrix} a_1^1 & a_2^1 & \cdots & a_n^1 \\ a_1^2 & a_2^2 & \cdots & a_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^m & a_2^m & \cdots & a_n^m \end{bmatrix} \begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} a_j^1 \\ a_j^2 \\ \vdots \\ a_j^m \end{bmatrix}$$

When

$$x = \sum_{j=1}^{n} \gamma^j x_j,$$

then

$$Ax = \sum_{j=1}^{n}\sum_{i=1}^{m} \gamma^j a_j^i y_i, = \sum_{i=1}^{m}\left(\sum_{j=1}^{n} \gamma^j a_j^i\right) y_i,$$

which gives rise to the familiar rule for matrix multiplication.

There is a one-to-one correspondence between matrices and linear operators in $L(X,Y)$. That is, once we fix a basis in X and in Y. If we would choose a different basis, we would get different matrices. This is important, the operator A acts on elements of X, the matrix is something that works with n-tuples of numbers.

Note that if B is an r-by-m matrix with entries b_k^j, then we note that the matrix for BA has the i,kth entry c_k^i being

$$c_k^i = \sum_{j=1}^{m} b_k^j a_j^i.$$

Note how upper and lower indices line up.

Now suppose that all the bases are just the standard bases and $X = \mathbb{R}^n$ and $Y = \mathbb{R}^m$. If we recall the Schwarz inequality we note that

$$\|Ax\|^2 = \sum_{i=1}^{m}\left(\sum_{j=1}^{n} \gamma^j a_j^i\right)^2 \leq \sum_{i=1}^{m}\left(\sum_{j=1}^{n} (\gamma^j)^2\right)\left(\sum_{j=1}^{n} (a_j^i)^2\right) = \sum_{i=1}^{m}\left(\sum_{j=1}^{n} (a_j^i)^2\right)\|x\|^2 .$$

In other words,

$$\|A\| \leq \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} (a_j^i)^2}.$$

Hence if the entries go to zero, then $\|A\|$ goes to zero. In particular, if the entries of $A-B$ go to zero then B goes to A in operator norm. That is in the metric space topology induced by the operator norm. We have proved the first part of:

Proposition: If $f\colon S \to \mathbb{R}^{nm}$ is a continuous function for a metric space S, then taking the components of f as the entries of a matrix, f is a continuous mapping from S to $L(\mathbb{R}^n,\mathbb{R}^m)$. Conversely if $f\colon S \to L(\mathbb{R}^n,\mathbb{R}^m)$ is a continuous function then the entries of the matrix are continuous functions.

The proof of the second part is rather easy. Take $f(x)e_j$ and note that is a continuous function to $\mathbb{R}^m$ with standard euclidean norm (Note $\|(A-B)e_j\| \leq \|A-B\|$). Such a function recall from last semester that such a function is continuous if and only if its components are continuous and these are the components of the jth column of the matrix $f(x)$.

The derivative

Recall that when we had a function $f\colon (a,b) \to \mathbb{R}$, we defined the derivative as

$$\lim_{h\to 0} \frac{f(x+h)-f(x)}{h}.$$

In other words, there was a number a such that

$$\lim_{h\to 0}\left|\frac{f(x+h)-f(x)}{h} - a\right| = \lim_{h\to 0}\left|\frac{f(x+h)-f(x)-ah}{h}\right| = \lim_{h\to 0}\frac{|f(x+h)-f(x)-ah|}{|h|} = 0.$$

Multiplying by a is a linear map in one dimension. That is, we can think of $a \in L(\mathbb{R}^1,\mathbb{R}^1)$. So we can use this definition to extend differentiation to more variables.

Definition: Let $U \subset \mathbb{R}^n$ be an open subset and $f\colon U \to \mathbb{R}^m$. Then we say f is *differentiable* at $x \in U$ if there exists an $A \in L(\mathbb{R}^n,\mathbb{R}^m)$ such that

$$\lim_{\substack{h\to 0\\ h\in\mathbb{R}^n}} \frac{\|f(x+h)-f(x)-Ah\|}{\|h\|} = 0.$$

and we say $Df(x) = A$, or $f'(x) = A$ and we say A is the *derivative* of f at x. When f is *differentiable* at all $x \in U$, we say simply that f is differentiable.

Note that the derivative is a function from U to $L(\mathbb{R}^n, \mathbb{R}^m)$.

The norms above must be in the right spaces of course. Normally it is just understood that $h \in \mathbb{R}^n$ and so we won't explicitly say so from now on.

We have again (as last semester) cheated somewhat as said that A is *the* derivative. We have of course not shown yet that there is only one.

Proposition: Let $U \subset \mathbb{R}^n$ be an open subset and $f\colon U \to \mathbb{R}^m$. Suppose that $x \in U$ and there exists an $A, B \in L(\mathbb{R}^n, \mathbb{R}^m)$ such that

$$\lim_{h\to 0} \frac{\|f(x+h) - f(x) - Ah\|}{\|h\|} = 0 \qquad \text{and} \qquad \lim_{h\to 0} \frac{\|f(x+h) - f(x) - Bh\|}{\|h\|} = 0.$$

Then $A = B$.

Proof.

$$\begin{aligned}\frac{\|(A-B)h\|}{\|h\|} &= \frac{\|f(x+h) - f(x) - Ah - (f(x+h) - f(x) - Bh)\|}{\|h\|} \\ &\leq \frac{\|f(x+h) - f(x) - Ah\|}{\|h\|} + \frac{\|f(x+h) - f(x) - Bh\|}{\|h\|}.\end{aligned}$$

So $\frac{\|(A-B)h\|}{\|h\|}$ goes to 0 as $h \to 0$. That is, given $\epsilon > 0$ we have that for all h in some δ ball around the origin

$$\epsilon > \frac{\|(A-B)h\|}{\|h\|} = \left\| (A-B)\frac{h}{\|h\|} \right\|.$$

For any x with $\|x\| = 1$ let $h = {}^{\delta}/_{2}\, x$, then $\|h\| < \delta$ and $\frac{h}{\|h\|} = x$ and so $\|A - B\| \leq \epsilon$. So $A = B$. $\square$

Example: If $f(x) = Ax$ for a linear mapping A, then $f'(x) = A$. This is easily seen:

$$\frac{\|f(x+h) - f(x) - Ah\|}{\|h\|} = \frac{\|A(x+h) - Ax - Ah\|}{\|h\|} = \frac{0}{\|h\|} = 0.$$

Proposition: Let $U \subset \mathbb{R}^n$ be open and $f\colon U \to \mathbb{R}^m$ be differentiable at x_0. Then f is continuous at x_0.

Proof. Another way to write the differentiability is to write

$$r(h) = f(x_0 + h) - f(x_0) - f'(x_0)h.$$

As $\frac{\|r(h)\|}{\|h\|}$ must go to zero as $h \to 0$, $r(h)$ itself must go to zero. As $h \mapsto f'(x_0)h$ is continuous (it's linear) and hence also goes to zero, that means that $f(x_0 + h)$ must go to $f(x_0)$. That is, f is continuous at x_0. $\square$

Theorem 9.15 (Chain rule): Let $U \subset \mathbb{R}^n$ be open and let $f\colon U \to \mathbb{R}^m$ be differentiable at $x_0 \in U$. Let $V \subset \mathbb{R}^m$ be open, $f(U) \subset V$ and let $g\colon V \to \mathbb{R}^\ell$ be differentiable at $f(x_0)$. Then

$$F(x) = g\big(f(x)\big)$$

is differentiable at x_0 and

$$F'(x_0) = g'\big(f(x_0)\big) f'(x_0).$$

Without the points this is sometimes written as $F' = g'f'$. The way to understand it is that the derivative of the composition $g\big(f(x)\big)$ is the composition of the derivatives of g and f. That is, if $A = f'(x_0)$ and $B = g'\big(f(x_0)\big)$, then $F'(x_0) = BA$.

Proof. Let $A = f'(x_0)$ and $B = g'\big(f(x_0)\big)$. Let h vary in $\mathbb{R}^n$ and write, $y_0 = f(x_0)$, $k = f(x_0 + h) - f(x_0)$. Write
$$r(h) = f(x_0 + h) - f(x_0) - Ah = k - Ah.$$
Then
$$\begin{aligned}
\frac{\|F(x_0+h) - F(x_0) - BAh\|}{\|h\|} &= \frac{\big\|g\big(f(x_0+h)\big) - g\big(f(x_0)\big) - BAh\big\|}{\|h\|} \\
&= \frac{\big\|g(y_0+k) - g(y_0) - B\big(k - r(h)\big)\big\|}{\|h\|} \\
&\leq \frac{\|g(y_0+k) - g(y_0) - Bk\|}{\|h\|} + \|B\| \frac{\|r(h)\|}{\|h\|} \\
&= \frac{\|g(y_0+k) - g(y_0) - Bk\|}{\|k\|} \frac{\|f(x_0+h) - f(x_0)\|}{\|h\|} + \|B\| \frac{\|r(h)\|}{\|h\|}.
\end{aligned}$$
First, $\|B\|$ is constant and f is differentiable at x_0, so the term $\|B\| \frac{\|r(h)\|}{\|h\|}$ goes to 0. Next as f is continuous at x_0, we have that as h goes to 0, then k goes to 0. Therefore $\frac{\|g(y_0+k)-g(y_0)-Bk\|}{\|k\|}$ goes to 0 because g is differentiable at y_0. Finally
$$\frac{\|f(x_0+h) - f(x_0)\|}{\|h\|} \leq \frac{\|f(x_0+h) - f(x_0) - Ah\|}{\|h\|} + \frac{\|Ah\|}{\|h\|} \leq \frac{\|f(x_0+h) - f(x_0) - Ah\|}{\|h\|} + \|A\|.$$
As f is differentiable at x_0, the term $\frac{\|f(x_0+h)-f(x_0)\|}{\|h\|}$ stays bounded as h goes to 0. Therefore, $\frac{\|F(x_0+h)-F(x_0)-BAh\|}{\|h\|}$ goes to zero, and hence $F'(x_0) = BA$, which is what was claimed. □

There is another way to generalize the derivative from one dimension. We can simply hold all but one variables constant and take the regular derivative.

Definition: Let $f \colon U \to \mathbb{R}$ be a function on an open set $U \subset \mathbb{R}^n$. If the following limit exists we write
$$\frac{\partial f}{\partial x^j}(x) = \lim_{h \to 0} \frac{f(x^1, \ldots, x^{j-1}, x^j + h, x^{j+1}, \ldots, x^n) - f(x)}{h} = \lim_{h \to 0} \frac{f(x + he_j) - f(x)}{h}.$$
We call $\frac{\partial f}{\partial x^j}(x)$ the *partial derivative* of f with respect to x^j. Sometimes we write $D_j f$ instead.

When $f \colon U \to \mathbb{R}^m$ is a function, then we can write $f = (f^1, f^2, \ldots, f^m)$, where f^k are real-valued functions. Then we can define $\frac{\partial f^k}{\partial x^j}$ (or write it $D_j f^k$).

Partial derivatives are easier to compute with all the machinery of calculus, and they provide a way to compute the total derivative of a function.

Theorem 9.17: Let $U \subset \mathbb{R}^n$ be open and let $f \colon U \to \mathbb{R}^m$ be differentiable at $x_0 \in U$. Then all the partial derivatives at x_0 exist and in terms of the standard basis of $\mathbb{R}^n$ and $\mathbb{R}^m$, $f'(x_0)$ is represented by the matrix
$$\begin{bmatrix}
\frac{\partial f^1}{\partial x^1}(x_0) & \frac{\partial f^1}{\partial x^2}(x_0) & \ldots & \frac{\partial f^1}{\partial x^n}(x_0) \\
\frac{\partial f^2}{\partial x^1}(x_0) & \frac{\partial f^2}{\partial x^2}(x_0) & \ldots & \frac{\partial f^2}{\partial x^n}(x_0) \\
\vdots & \vdots & \ddots & \vdots \\
\frac{\partial f^m}{\partial x^1}(x_0) & \frac{\partial f^m}{\partial x^2}(x_0) & \ldots & \frac{\partial f^m}{\partial x^n}(x_0)
\end{bmatrix}.$$

In other words
$$f'(x_0)\, e_j = \sum_{k=1}^{m} \frac{\partial f^k}{\partial x^j}(x_0)\, e_k.$$
If $h = \sum_{j=1}^{n} c^j e_j$, then
$$f'(x_0)\, h = \sum_{j=1}^{n} \sum_{k=1}^{m} c^j \frac{\partial f^k}{\partial x^j}(x_0)\, e_k.$$
Again note the up-down pattern with the indices being summed over. That is on purpose.

Proof. Fix a j and note that

$$\left\| \frac{f(x_0+he_j)-f(x_0)}{h} - f'(x_0)e_j \right\| = \left\| \frac{f(x_0+he_j)-f(x_0)-f'(x_0)he_j}{h} \right\| = \frac{\|f(x_0+he_j)-f(x_0)-f'(x_0)he_j\|}{\|he_j\|}.$$

As h goes to 0, the right hand side goes to zero by differentiability of f, and hence

$$\lim_{h\to 0} \frac{f(x_0+he_j)-f(x_0)}{h} = f'(x_0)e_j.$$

Note that f is vector valued. So represent f by components $f = (f^1, f^2, \dots, f^m)$, and note that taking a limit in $\mathbb{R}^m$ is the same as taking the limit in each component separately. Therefore for any k the partial derivative

$$\frac{\partial f^k}{\partial x^j}(x_0) = \lim_{h\to 0} \frac{f^k(x_0+he_j)-f^k(x_0)}{h}$$

exists and is equal to the kth component of $f'(x_0)e_j$, and we are done. $\square$

One of the consequences of the theorem is that if f is differentiable on U, then $f' \colon U \to L(\mathbb{R}^n, \mathbb{R}^m)$ is a continuous function if and only if all the $\frac{\partial f^k}{\partial x^j}$ are continuous functions.

When $U \subset \mathbb{R}^n$ is open and $f \colon U \to \mathbb{R}$ is a differentiable function. We call the following vector the *gradient*:

$$\nabla f(x) = \sum_{j=1}^n \frac{\partial f}{\partial x^j}(x)\, e_j.$$

Note that the upper-lower indices don't really match up. As a preview of Math 621, we note that we write

$$df = \sum_{j=1}^n \frac{\partial f}{\partial x^j} dx^j$$

where dx^j is really the standard bases (though we're thinking of dx^j to be in $L(\mathbb{R}^n, \mathbb{R})$ which is really equivalent to $\mathbb{R}^n$). But we digress.

Suppose that $\gamma \colon (a,b) \subset \mathbb{R} \to \mathbb{R}^n$ is a differentiable function and $\gamma\big((a,b)\big) \subset U$. Write $\gamma = (\gamma^1, \dots, \gamma^n)$. Then we can write

$$g(t) = f\big(\gamma(t)\big).$$

The function g is then a differentiable and the derivative is

$$g'(t) = \sum_{j=1}^n \frac{\partial f}{\partial x^j}\big(\gamma(t)\big) \frac{d\gamma^j}{dt}(t) = \sum_{j=1}^n \frac{\partial f}{\partial x^j} \frac{d\gamma^j}{dt},$$

where we sometimes, for convenience, leave out the points at which we are evaluating. We notice that

$$g'(t) = (\nabla f)\big(\gamma(t)\big) \cdot \gamma'(t) = \nabla f \cdot \gamma'.$$

The dot represents the standard scalar dot product.

We use this idea to define derivatives in a specific direction. A direction is simply a vector pointing in that direction. So pick a vector $u \in \mathbb{R}^n$ such that $\|u\| = 1$. Fix $x \in U$. Then define

$$\gamma(t) = x + tu.$$

It is easy to compute that $\gamma'(t) = u$ for all t. By chain rule

$$\frac{d}{dt}\Big|_{t=0} \big[f(x+tu)\big] = (\nabla f)(x) \cdot u,$$

where the notation $\frac{d}{dt}|_{t=0}$ represents the derivative evaluated at $t = 0$. We also just compute directly

$$\frac{d}{dt}\Big|_{t=0} \big[f(x+tu)\big] = \lim_{h\to 0} \frac{f(x+hu)-f(x)}{h}.$$

We obtain what is usually called the *directional derivative*, sometimes denoted by

$$D_u f(x) = \frac{d}{dt}\Big|_{t=0} \big[f(x + tu) \big],$$

which can be computed by one of the methods above.

Let us suppose that $(\nabla f)(x) \neq 0$. By Schwarz inequality we have

$$\left| D_u f(x) \right| \leq \left\| (\nabla f)(x) \right\|,$$

and further equality is achieved when u is a scalar multiple of $(\nabla f)(x)$. When

$$u = \frac{(\nabla f)(x)}{\left\| (\nabla f)(x) \right\|},$$

we get $D_u f(x) = \left\| (\nabla f)(x) \right\|$. So the gradient points in the direction in which the function grows fastest, that is, the direction in which D_u is maximal.

Let us prove a "mean value theorem" for vector valued functions.

Theorem 5.19: If $\varphi \colon [a, b] \to \mathbb{R}^n$ is differentiable on (a, b) and continuous on $[a, b]$, then there exists a t such that

$$\left\| \varphi(b) - \varphi(a) \right\| \leq (b - a) \left\| \varphi'(t) \right\|.$$

Proof. By mean value theorem on the function $\big(\varphi(b) - \varphi(a)\big) \cdot \varphi(t)$ (the dot is the scalar dot product) we obtain there is a t such that

$$\big(\varphi(b) - \varphi(a)\big) \cdot \varphi(b) - \big(\varphi(b) - \varphi(a)\big) \cdot \varphi(a) = \left\| \varphi(b) - \varphi(a) \right\|^2 = \big(\varphi(b) - \varphi(a)\big) \cdot \varphi'(t)$$

where we treat φ' as a simply a column vector of numbers by abuse of notation. Note that in this case, it is not hard to see that $\left\| \varphi'(t) \right\|_{L(\mathbb{R}, \mathbb{R}^n)} = \left\| \varphi'(t) \right\|_{\mathbb{R}^n}$ (exercise).

By Schwarz inequality

$$\left\| \varphi(b) - \varphi(a) \right\|^2 = \big(\varphi(b) - \varphi(a)\big) \cdot \varphi'(t) \leq \left\| \varphi(b) - \varphi(a) \right\| \left\| \varphi'(t) \right\|.$$

$\square$

A set U is *convex* if whenever $x, y \in U$, the line segment from x to y lies in U. That is, if the *convex combination* $(1 - t)x + ty$ is in U for all $t \in [0, 1]$. Note that in $\mathbb{R}$, every connected interval is convex. In $\mathbb{R}^2$ (or higher dimensions) there are lots of nonconvex connected sets. The ball $B(x, r)$ is always convex by the triangle inequality (exercise).

Theorem 9.19: Let $U \subset \mathbb{R}^n$ be a convex open set, $f \colon U \to \mathbb{R}^m$ a differentiable function, and an M such that

$$\left\| f'(x) \right\| \leq M$$

for all $x \in U$. Then f is Lipschitz with constant M, that is

$$\left\| f(x) - f(y) \right\| \leq M \left\| x - y \right\|$$

for all $x, y \in U$.

Note that if U is not convex this is not true. To see this, take the set

$$U = \{(x, y) : 0.9 < x^2 + y^2 < 1.1\} \setminus \{(x, 0) : x < 0\}.$$

Let $f(x, y)$ be the angle that the line from the origin to (x, y) makes with the positive x axis. You can even write the formula for f:

$$f(x, y) = 2 \arctan \left(\frac{y}{x + \sqrt{x^2 + y^2}} \right).$$

Think spiral staircase with room in the middle. See:

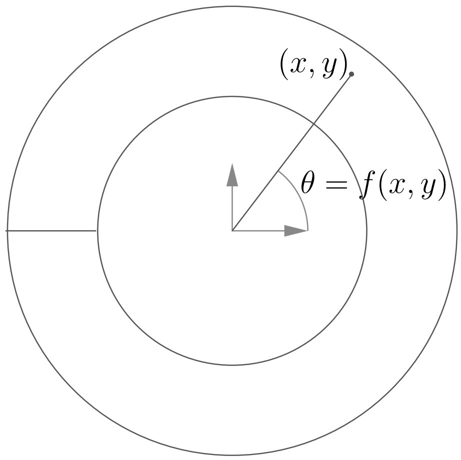

In any case the function is differentiable, and the derivative is bounded on U, which is not hard to see, but thinking of what happens near where the negative x-axis cuts the annulus in half, we see that the conclusion cannot hold.

Proof. Fix x and y in U and note that $(1-t)x + ty \in U$ for all $t \in [0,1]$ by convexity. Next

$$\frac{d}{dt}\Big[f\big((1-t)x + ty\big)\Big] = f'\big((1-t)x + ty\big)(y - x).$$

By mean value theorem above we get

$$\|f(x) - f(y)\| \leq \left\| \frac{d}{dt}\Big[f\big((1-t)x + ty\big)\Big] \right\| \leq \left\|f'\big((1-t)x + ty\big)\right\| \|y - x\| \leq M \|y - x\| .$$

$\square$

Let us solve the differential equation $f' = 0$.

Corollary: If $U \subset \mathbb{R}^n$ is connected and $f \colon U \to \mathbb{R}^m$ is differentiable and $f'(x) = 0$, for all $x \in U$, then f is constant.

Proof. For any $x \in U$, there is a ball $B(x, \delta) \subset U$. The ball $B(x, \delta)$ is convex. Since $\|f'(y)\| \leq 0$ for all $y \in B(x, \delta)$ then by the theorem, $\|f(x) - f(y)\| \leq 0 \|x - y\| = 0$, so $f(x) = f(y)$ for all $y \in B(x, \delta)$.

This means that $f^{-1}(c)$ is open for any $c \in \mathbb{R}^m$. Suppose that $f^{-1}(c)$ is nonempty. The two sets

$$U' = f^{-1}(c), \qquad U'' = f^{-1}(\mathbb{R}^m \setminus \{c\}) = \bigcup_{\substack{a \in \mathbb{R}^m \\ a \neq c}} f^{-1}(a)$$

are open disjoint, and further $U = U' \cup U''$. So as U' is nonempty, and U is connected, we have that $U'' = \emptyset$. So $f(x) = c$ for all $x \in U$. $\square$

Definition: We say $f \colon U \subset \mathbb{R}^n \to \mathbb{R}^m$ is *continuously differentiable*, or $C^1(U)$ if f is differentiable and $f' \colon U \to L(\mathbb{R}^n, \mathbb{R}^m)$ is continuous.

Theorem 9.21: Let $U \subset \mathbb{R}^n$ be open and $f \colon U \to \mathbb{R}^m$. The function f is continuously differentiable if and only if all the partial derivatives exist and are continuous.

Note that without continuity the theorem does not hold. Just because partial derivatives exist doesn't mean that f is differentiable, in fact, f may not even be continuous. See the homework.

Proof. We have seen that if f is differentiable, then the partial derivatives exist. Furthermore, the partial derivatives are the entries of the matrix of $f'(x)$. So if $f' \colon U \to L(\mathbb{R}^n, \mathbb{R}^m)$ is continuous, then the entries are continuous, hence the partial derivatives are continuous.

To prove the opposite direction, suppose the partial derivatives exist and are continuous. Fix $x \in U$. If we can show that $f'(x)$ exists we are done, because the entries of the matrix $f'(x)$ are then the partial derivatives and if the entries are continuous functions, the matrix valued function f' is continuous.

Let us do induction on dimension. First let us note that the conclusion is true when $n = 1$. In this case the derivative is just the regular derivative (exercise: you should check that the fact that the function is vector valued is not a problem).

Suppose that the conclusion is true for $\mathbb{R}^{n-1}$, that is, if we restrict to the first $n-1$ variables, the conclusion is true. It is easy to see that the first $n-1$ partial derivatives of f restricted to the set where the last coordinate is fixed are the same as those for f. In the following we will think of $\mathbb{R}^{n-1}$ as a subset of $\mathbb{R}^n$, that is the set in $\mathbb{R}^n$ where $x^n = 0$. Let

$$A = \begin{bmatrix} \frac{\partial f^1}{\partial x^1}(x) & \dots & \frac{\partial f^1}{\partial x^n}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial f^m}{\partial x^1}(x) & \dots & \frac{\partial f^m}{\partial x^n}(x) \end{bmatrix}, \qquad A_1 = \begin{bmatrix} \frac{\partial f^1}{\partial x^1}(x) & \dots & \frac{\partial f^1}{\partial x^{n-1}}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial f^m}{\partial x^1}(x) & \dots & \frac{\partial f^m}{\partial x^{n-1}}(x) \end{bmatrix}, \qquad v = \begin{bmatrix} \frac{\partial f^1}{\partial x^n}(x) \\ \vdots \\ \frac{\partial f^m}{\partial x^n}(x) \end{bmatrix}.$$

Let $\epsilon > 0$ be given. Let $\delta > 0$ be such that for any $k \in \mathbb{R}^{n-1}$ with $\|k\| < \delta$ we have

$$\frac{\|f(x+k) - f(x) - A_1 k\|}{\|k\|} < \epsilon.$$

By continuity of the partial derivatives, suppose that δ is small enough so that

$$\left| \frac{\partial f^j}{\partial x^n}(x+h) - \frac{\partial f^j}{\partial x^n}(x) \right| < \epsilon,$$

for all j and all h with $\|h\| < \delta$.

Let $h = h_1 + t e_n$ be a vector in $\mathbb{R}^n$ where $h_1 \in \mathbb{R}^{n-1}$ such that $\|h\| < \delta$. Then $\|h_1\| \leq \|h\| < \delta$. Note that $Ah = A_1 h_1 + tv$.

$$\begin{aligned} \|f(x+h) - f(x) - Ah\| &= \|f(x+h_1+te_n) - f(x+h_1) - tv + f(x+h_1) - f(x) - A_1 h_1\| \\ &\leq \|f(x+h_1+te_n) - f(x+h_1) - tv\| + \|f(x+h_1) - f(x) - A_1 h_1\| \\ &\leq \|f(x+h_1+te_n) - f(x+h_1) - tv\| + \epsilon \|h_1\|. \end{aligned}$$

As all the partial derivatives exist then by the mean value theorem for each j there is some $\theta_j \in [0,t]$ (or $[t,0]$ if $t < 0$), such that

$$f^j(x+h_1+te_n) - f^j(x+h_1) = t\frac{\partial f^j}{\partial x^n}(x+h_1+\theta_j e_n).$$

Note that if $\|h\| < \delta$ then $\|h_1 + \theta_j e_n\| \leq \|h\| < \delta$. So to finish the estimate

$$\begin{aligned} \|f(x+h) - f(x) - Ah\| &\leq \|f(x+h_1+te_n) - f(x+h_1) - tv\| + \epsilon \|h_1\| \\ &\leq \sqrt{\sum_{j=1}^{m} \left(t\frac{\partial f^j}{\partial x^n}(x+h_1+\theta_j e_n) - t\frac{\partial f^j}{\partial x^n}(x) \right)^2} + \epsilon \|h_1\| \\ &\leq \sqrt{m}\, \epsilon \, |t| + \epsilon \|h_1\| \\ &\leq (\sqrt{m}+1)\epsilon \|h\|. \end{aligned}$$

□

Contraction mapping principle

Let us review the contraction mapping principle.

Definition: Let (X,d) and (X',d') be metric spaces. $f \colon X \to X'$ is said to be a *contraction* (or a contractive map) if it is a k-Lipschitz map for some $k < 1$, i.e. if there exists a $k < 1$ such that

$$d'\big(f(x), f(y)\big) \leq k d(x,y) \quad \text{for all } x, y \in X.$$

If $f \colon X \to X$ is a map, $x \in X$ is called a *fixed point* if $f(x) = x$.

Theorem 9.23 (Contraction mapping principle or Fixed point theorem): Let (X,d) be a nonempty complete metric space and $f \colon X \to X$ a contraction. Then f has a fixed point.

The words *complete* and *contraction* are necessary. For example, $f\colon (0,1) \to (0,1)$ defined by $f(x) = kx$ for any $0 < k < 1$ is a contraction with no fixed point. Also $f\colon \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x + 1$ is not a contraction ($k = 1$) and has no fixed point.

Proof. Pick any $x_0 \in X$. Define a sequence $\{x_n\}$ by $x_{n+1} := f(x_n)$.

$$d(x_{n+1}, x_n) = d\big(f(x_n), f(x_{n-1})\big) \leq kd(x_n, x_{n-1}) \leq \cdots \leq k^n d(x_1, x_0).$$

Suppose $m \geq n$, then

$$\begin{aligned} d(x_m, x_n) &\leq \sum_{i=n}^{m-1} d(x_{i+1}, x_i) \\ &\leq \sum_{i=n}^{m-1} k^i d(x_1, x_0) \\ &= k^n d(x_1, x_0) \sum_{i=0}^{m-n-1} k^i \\ &\leq k^n d(x_1, x_0) \sum_{i=0}^{\infty} k^i = k^n d(x_1, x_0) \frac{1}{1-k}. \end{aligned}$$

In particular the sequence is Cauchy (why?). Since X is complete we let $x := \lim\, x_n$ and we claim that x is our unique fixed point.

Fixed point? Note that f is continuous because it is a contraction. Hence

$$f(x) = \lim\, f(x_n) = \lim\, x_{n+1} = x.$$

Unique? Let y be a fixed point.

$$d(x, y) = d\big(f(x), f(y)\big) = kd(x, y).$$

As $k < 1$ this means that $d(x, y) = 0$ and hence $x = y$. The theorem is proved. □

Note that the proof is constructive. Not only do we know that a unique fixed point exists. We also know how to find it.

We've used the theorem to prove Picard's theorem last semester. This semester, we will prove the inverse and implicit function theorems.

Do also note the proof of uniqueness holds even if X is not complete. If f is a contraction, then if it has a fixed point, that point is unique.

Inverse function theorem

The idea of a derivative is that if a function is differentiable, then it locally "behaves like" the derivative (which is a linear function). So for example, if a function is differentiable and the derivative is invertible, the function is (locally) invertible.

Theorem 9.24: Let $U \subset \mathbb{R}^n$ be a set and let $f\colon U \to \mathbb{R}^n$ be a continuously differentiable function. Also suppose that $x_0 \in U$, $f(x_0) = y_0$, and $f'(x_0)$ is invertible. Then there exist open sets $V, W \subset \mathbb{R}^n$ such that $x_0 \in V \subset U$, $f(V) = W$ and $f|_V$ is one-to-one and onto. Furthermore, the inverse $g(y) = (f|_V)^{-1}(y)$ is continuously differentiable and

$$g'(y) = \big(f'(x)\big)^{-1}, \qquad \text{for all } x \in V,\ y = f(x).$$

Proof. Write $A = f'(x_0)$. As f' is continuous, there exists an open ball V around x_0 such that

$$\|A - f'(x)\| < \frac{1}{2\,\|A^{-1}\|} \qquad \text{for all } x \in V.$$

Note that $f'(x)$ is invertible for all $x \in V$.

Given $y \in \mathbb{R}^n$ we define $\varphi_y \colon C \to \mathbb{R}^n$

$$\varphi_y(x) = x + A^{-1}\big(y - f(x)\big).$$

As A^{-1} is one-to-one, we notice that $\varphi_y(x) = x$ (x is a fixed point) if only if $y - f(x) = 0$, or in other words $f(x) = y$. Using chain rule we obtain.

$$\varphi_y'(x) = I - A^{-1}f'(x) = A^{-1}\big(A - f'(x)\big).$$

so for $x \in V$ we have

$$\left\|\varphi_y'(x)\right\| \leq \left\|A^{-1}\right\| \left\|A - f'(x)\right\| < 1/2.$$

As V is a ball it is convex, and hence

$$\left\|\varphi_y(x_1) - \varphi_y(x_2)\right\| \leq \frac{1}{2}\left\|x_1 - x_2\right\| \qquad \text{for all } x_1, x_2 \in V.$$

In other words φ_y is a contraction defined on V, though we so far do not know what is the range of φ_y. We cannot apply the fixed point theorem, but we can say that φ_y has at most one fixed point (note proof of uniqueness in the contraction mapping principle). That is, there exists at most one $x \in V$ such that $f(x) = y$, and so $f|_V$ is one-to-one.

Let $W = f(V)$. We need to show that W is open. Take a $y_1 \in W$, then there is a unique $x_1 \in V$ such that $f(x_1) = y_1$. Let $r > 0$ be small enough such that the closed ball $C(x_1, r) \subset V$ (such $r > 0$ exists as V is open).

Suppose y is such that

$$\left\|y - y_1\right\| < \frac{r}{2\left\|A^{-1}\right\|}.$$

If we can show that $y \in W$, then we have shown that W is open. Define $\varphi_y(x) = x + A^{-1}\big(y - f(x)\big)$ as before. If $x \in C(x_1, r)$, then

$$\begin{aligned}
\left\|\varphi_y(x) - x_1\right\| &\leq \left\|\varphi_y(x) - \varphi_y(x_1)\right\| + \left\|\varphi_y(x_1) - x_1\right\| \\
&\leq \frac{1}{2}\left\|x - x_1\right\| + \left\|A^{-1}(y - y_1)\right\| \\
&\leq \frac{1}{2}r + \left\|A^{-1}\right\| \left\|y - y_1\right\| \\
&< \frac{1}{2}r + \left\|A^{-1}\right\| \frac{r}{2\left\|A^{-1}\right\|} = r.
\end{aligned}$$

So φ_y takes $C(x_1, r)$ into $B(x_1, r) \subset C(x_1, r)$. It is a contraction on $C(x_1, r)$ and $C(x_1, r)$ is complete (closed subset of $\mathbb{R}^n$ is complete). Apply the contraction mapping principle to obtain a fixed point x, i.e. $\varphi_y(x) = x$. That is $f(x) = y$. So $y \in f\big(C(x_1, r)\big) \subset f(V) = W$. Therefore W is open.

Next we need to show that g is continuously differentiable and compute its derivative. First let us show that it is differentiable. Let $y \in W$ and $k \in \mathbb{R}^n$, $k \neq 0$, such that $y + k \in W$. Then there are unique $x \in V$ and $h \in \mathbb{R}^n$, $h \neq 0$ and $x + h \in V$, such that $f(x) = y$ and $f(x + h) = y + k$ as $f|_V$ is a one-to-one and onto mapping of V onto W. In other words, $g(y) = x$ and $g(y + k) = x + h$. We can still squeeze some information from the fact that φ_y is a contraction.

$$\varphi_y(x + h) - \varphi_y(x) = h + A^{-1}\big(f(x) - f(x + h)\big) = h - A^{-1}k.$$

So

$$\left\|h - A^{-1}k\right\| = \left\|\varphi_y(x + h) - \varphi_y(x)\right\| \leq \frac{1}{2}\left\|x + h - x\right\| = \frac{\left\|h\right\|}{2}.$$

By the inverse triangle inequality $\left\|h\right\| - \left\|A^{-1}k\right\| \leq \frac{1}{2}\left\|h\right\|$ so

$$\left\|h\right\| \leq 2\left\|A^{-1}k\right\| \leq 2\left\|A^{-1}\right\| \left\|k\right\|.$$

In particular as k goes to 0, so does h.

As $x \in V$, then $f'(x)$ is invertible. Let $B = \big(f'(x)\big)^{-1}$, which is what we think the derivative of g at y is. Then

$$\begin{aligned}
\frac{\|g(y+k)-g(y)-Bk\|}{\|k\|} &= \frac{\|h-Bk\|}{\|k\|} \\
&= \frac{\big\|h-B\big(f(x+h)-f(x)\big)\big\|}{\|k\|} \\
&= \frac{\big\|B\big(f(x+h)-f(x)-f'(x)h\big)\big\|}{\|k\|} \\
&\leq \|B\| \frac{\|h\|}{\|k\|} \frac{\|f(x+h)-f(x)-f'(x)h\|}{\|h\|} \\
&\leq 2\,\|B\|\,\|A^{-1}\| \frac{\|f(x+h)-f(x)-f'(x)h\|}{\|h\|}.
\end{aligned}$$

As k goes to 0, so does h. So the right hand side goes to 0 as f is differentiable, and hence the left hand side also goes to 0. And B is precisely what we wanted $g'(y)$ to be.

We have that g is differentiable, let us show it is $C^1(W)$. Now, $g\colon W \to V$ is continuous (it's differentiable), f' is continuous function from V to $L(\mathbb{R}^n)$, and $X \to X^{-1}$ is a continuous function. $g'(y) = \Big(f'\big(g(y)\big)\Big)^{-1}$ is the composition of these three continuous functions and hence is continuous. □

Corollary: Suppose $U \subset \mathbb{R}^n$ is open and $f\colon U \to \mathbb{R}^n$ is a continuously differentiable mapping such that $f'(x)$ is invertible for all $x \in U$. Then given any open set $V \subset U$, $f(V)$ is open. (f is an open mapping).

Proof. WLOG suppose $U = V$. For each point $y \in f(V)$, we pick $x \in f^{-1}(y)$ (there could be more than one such point), then by the inverse function theorem there is a neighbourhood of x in V that maps onto an neighbourhood of y. Hence $f(V)$ is open. □

The theorem, and the corollary, is not true if $f'(x)$ is not invertible for some x. For example, the map $f(x,y) = (x,xy)$, maps $\mathbb{R}^2$ onto the set $\mathbb{R}^2 \setminus \{(0,y) : y \neq 0\}$, which is neither open nor closed. In fact $f^{-1}(0,0) = \{(0,y) : y \in \mathbb{R}\}$. Note that this bad behaviour only occurs on the y-axis, everywhere else the function is locally invertible. In fact if we avoid the y-axis it is even one to one.

Also note that just because $f'(x)$ is invertible everywhere doesn't mean that f is one-to-one globally. It is definitely "locally" one-to-one. For an example, just take the map $f\colon \mathbb{C} \setminus \{0\} \to \mathbb{C}$ defined by $f(z) = z^2$. Here we treat the map as if it went from $\mathbb{R}^2 \setminus \{0\}$ to $\mathbb{R}^2$. For any nonzero complex number, there are always two square roots, so the map is actually 2-to-1. It is left to student to show that f is differentiable and the derivative is invertible (Hint: let $z = x + iy$ and write down what the real an imaginary part of f is in terms if x and y).

Also note that the invertibility of the derivative is not a necessary condition, just sufficient for having a continuous inverse and being an open mapping. For example the function $f(x) = x^3$ is an open mapping from $\mathbb{R}$ to $\mathbb{R}$ and is globally one-to-one with a continuous inverse.

Implicit function theorem:

The inverse function theorem is really a special case of the implicit function theorem which we prove next. Although somewhat ironically we will prove the implicit function theorem using the inverse function theorem. Really what we were showing in the inverse function theorem was that the equation $x - f(y) = 0$ was solvable for y in terms of x if the derivative in terms of y was invertible, that is if $f'(y)$ was invertible. That is there was locally a function g such that $x - f\big(g(x)\big) = 0$.

OK, so how about we look at the equation $f(x,y) = 0$. Obviously this is not solvable for y in terms of x in every case. For example, when $f(x,y)$ does not actually depend on y. For a slightly more complicated example, notice that $x^2 + y^2 - 1 = 0$ defines the unit circle, and we can locally solve for y in terms of x when 1) we are near a point which lies on the unit circle and 2) when we are not at a point where the circle has a vertical tangency, or in other words where $\frac{\partial f}{\partial y} = 0$.

To make things simple we fix some notation. We let $(x, y) \in \mathbb{R}^{n+m}$ denote the coordinates $(x^1, \ldots, x^n, y^1, \ldots, y^m)$ A linear transformation $A \in L(\mathbb{R}^{n+m}, \mathbb{R}^m)$ can then always be written as $A = [A_x \; A_y]$ so that $A(x, y) = A_x x + A_y y$, where $A_x \in L(\mathbb{R}^n, \mathbb{R}^m)$ and $A_y \in L(\mathbb{R}^m)$.

Note that Rudin does things "in reverse" from what the statement is usually. I'll do it in the usual order as that's what I am used to, where we are taking the derivatives of y, not x (but it doesn't matter really in the end). First a linear version of the implicit function theorem.

Proposition (Theorem 9.27): Let $A = [A_x \; A_y] \in L(\mathbb{R}^{n+m}, \mathbb{R}^m)$ and suppose that A_y is invertible, then let $B = -(A_y)^{-1} A_x$ and note that

$$0 = A(x, Bx) = A_x x + A_y Bx.$$

The proof is obvious. We simply solve and obtain $y = Bx$. Let us therefore show that the same can be done for C^1 functions.

Theorem 9.28 (Implicit function theorem): Let $U \subset \mathbb{R}^{n+m}$ be an open set and let $f \colon U \to \mathbb{R}^m$ be a $C^1(U)$ mapping. Let $(x_0, y_0) \in U$ be a point such that $f(x_0, y_0) = 0$. Write $A = [A_x \; A_y] = f'(x_0, y_0)$ and suppose that A_y is invertible. Then there exists an open set $W \subset \mathbb{R}^n$ with $x_0 \in W$ and a $C^1(W)$ mapping $g \colon W \to \mathbb{R}^m$, with $g(x_0) = y_0$, and for all $x \in W$, we have $(x, g(x)) \in U$ and

$$f\bigl(x, g(x)\bigr) = 0.$$

Furthermore,

$$g'(x_0) = -(A_y)^{-1} A_x.$$

Proof. Define $F \colon U \to \mathbb{R}^{n+m}$ by $F(x, y) = \bigl(x, f(x, y)\bigr)$. It is clear that F is C^1, and we want to show that the derivative at (x_0, y_0) is invertible.

Let's compute the derivative. We know that

$$\frac{\|f(x_0 + h, y_0 + k) - f(x_0, y_0) - A_x h - A_y k\|}{\|(h, k)\|}$$

goes to zero as $\|(h, k)\| = \sqrt{\|h\|^2 + \|k\|^2}$ goes to zero. But then so does

$$\frac{\bigl\|\bigl(h, f(x_0 + h, y_0 + k) - f(x_0, y_0)\bigr) - (h, A_x h + A_y k)\bigr\|}{\|(h, k)\|} = \frac{\|f(x_0 + h, y_0 + k) - f(x_0, y_0) - A_x h - A_y k\|}{\|(h, k)\|}.$$

So the derivative of F ate (x_0, y_0) takes (h, k) to $(h, A_x h + A_y k)$. If $(h, A_x h + A_y k) = (0, 0)$, then $h = 0$, and so $A_y k = 0$. As A_y is one-to-one, then $k = 0$. Therefore $F'(x_0, y_0)$ is one-to-one or in other words invertible and we can apply the inverse function theorem.

That is, there exists some open set $V \subset \mathbb{R}^{n+m}$ with $(x_0, 0) \in V$, and an inverse mapping $G \colon V \to \mathbb{R}^{n+m}$, that is $F(G(x, s)) = (x, s)$ for all $(x, s) \in V$ (where $x \in \mathbb{R}^n$ and $s \in \mathbb{R}^m$). Write $G = (G_1, G_2)$ (the first n and the second m components of G). Then

$$F\bigl(G_1(x, s), G_2(x, s)\bigr) = \bigl(G_1(x, s), f(G_1(x, s), G_2(x, s))\bigr) = (x, s).$$

So $x = G_1(x, s)$ and $f\bigl(G_1(x, s), G_2(x, s)\bigr) = f\bigl(x, G_2(x, s)\bigr) = s$. Plugging in $s = 0$ we obtain

$$f\bigl(x, G_2(x, 0)\bigr) = 0.$$

Let $W = \{x \in \mathbb{R}^n : (x, 0) \in V\}$ and define $g \colon W \to \mathbb{R}^m$ by $g(x) = G_2(x, 0)$. We obtain the g in the theorem.

Next differentiate

$$x \mapsto f\bigl(x, g(x)\bigr),$$

at x_0, which should be the zero map. The derivative is done in the same way as above. We get that for all $h \in \mathbb{R}^n$

$$0 = A\bigl(h, g'(x_0)h\bigr) = A_x h + A_y g'(x_0) h,$$

and we obtain the desired derivative for g as well. $\square$

In other words, in the context of the theorem we have m equations in $n+m$ unknowns.

$$\begin{gathered} f^1(x_1,\ldots,x_n,y_1,\ldots,y_m)=0 \\ \vdots \\ f^m(x_1,\ldots,x_n,y_1,\ldots,y_m)=0 \end{gathered}$$

And the condition guaranteeing a solution is that this is a C^1 mapping (that all the components are C^1, or in other words all the partial derivatives exist and are continuous), and the matrix

$$\begin{bmatrix} \frac{\partial f^1}{\partial y^1} & \cdots & \frac{\partial f^1}{\partial y^m} \\ \vdots & \ddots & \vdots \\ \frac{\partial f^m}{\partial y^1} & \cdots & \frac{\partial f^m}{\partial y^m} \end{bmatrix}$$

is invertible at (x_0, y_0).

Example: Consider the set $x^2+y^2-(z+1)^3=-1$, $e^x+e^y+e^z=3$ near the point $(0,0,0)$. The function we are looking at is

$$f(x,y,z)=(x^2+y^2-(z+1)^3+1, e^x+e^y+e^z-3).$$

We find that

$$Df=\begin{bmatrix} 2x & 2y & -3(z+1)^2 \\ e^x & e^y & e^z \end{bmatrix}.$$

The matrix

$$\begin{bmatrix} 2(0) & -3(0+1)^2 \\ e^0 & e^0 \end{bmatrix}=\begin{bmatrix} 0 & -3 \\ 1 & 1 \end{bmatrix}$$

is invertible. Hence near $(0,0,0)$ we can find y and z as C^1 functions of x such that for x near 0 we have

$$x^2+y(x)^2-(z(x)+1)^3=-1, \qquad e^x+e^{y(x)}+e^{z(x)}=3.$$

The theorem doesn't tell us how to find $y(x)$ and $z(x)$ explicitly, it just tells us they exist. In other words, near the origin the set of solutions is a smooth curve that goes through the origin.

Note that there are versions of the theorem for arbitrarily many derivatives. If f has k continuous derivatives, then the solution also has k derivatives.

So it would be good to have an easy test for when is a matrix invertible. This is where determinants come in. Suppose that $\sigma=(\sigma_1,\ldots,\sigma_n)$ is a permutation of the integers $(1,\ldots,n)$. It is not hard to see that any permutation can be obtained by a sequence of transpositions (switchings of two elements). Call a permutation even (resp. odd) if it takes an even (resp. odd) number of transpositions to get from σ to $(1,\ldots,n)$. It can be shown that this is well defined, in fact it is not hard to show that

$$\operatorname{sgn}(\sigma)=\operatorname{sgn}(\sigma_1,\ldots,\sigma_n)=\prod_{p<q}\operatorname{sgn}(\sigma_q-\sigma_p)$$

is -1 if σ is odd and 1 if σ is even. The symbol $\operatorname{sgn}(x)$ for a number is defined by

$$\operatorname{sgn}(x)=\begin{cases} -1 & \text{if } x<0, \\ 0 & \text{if } x=0, \\ 1 & \text{if } x>0. \end{cases}$$

This can be proved by noting that applying a transposition changes the sign, which is not hard to prove by induction on n. Then note that the sign of $(1,2,\ldots,n)$ is 1.

Let S_n be the set of all permutations on n elements (the *symmetric group*). Let $A=[a_j^i]$ be a matrix. Define the *determinant* of A

$$\det(A)=\sum_{\sigma\in S_n}\operatorname{sgn}(\sigma)\prod_{i=1}^n a_{\sigma_i}^i.$$

Proposition (Theorem 9.34 and other observations):

(i) $\det(I) = 1$.
(ii) $\det([x_1 x_2 \ldots x_n])$ where x_j are column vectors is linear in each variable x_j separately.
(iii) If two columns of a matrix are interchanged determinant changes sign.
(iv) If two columns of A are equal, then $\det(A) = 0$.
(v) If a column is zero, then $\det(A) = 0$.
(vi) $A \mapsto \det(A)$ is a continuous function.
(vii) $\det \left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right] = ad - bc$ and $\det[a] = a$.

In fact, the determinant is the unique function that satisfies (i), (ii), and (iii). But we digress.

Proof. We go through the proof quickly, as you have likely seen this before.

(i) is trivial. For (ii) Notice that each term in the definition of the determinant contains exactly one factor from each column.

Part (iii) follows by noting that switching two columns is like switching the two corresponding numbers in every element in S_n. Hence all the signs are changed. Part (iv) follows because if two columns are equal and we switch them we get the same matrix back and so part (iii) says the determinant must have been 0.

Part (v) follows because the product in each term in the definition includes one element from the zero column. Part (vi) follows as det is a polynomial in the entries of the matrix and hence continuous. We have seen that a function defined on matrices is continuous in the operator norm if it is continuous in the entries. Finally, part (vii) is a direct computation. $\square$

Theorem 9.35+9.36: If A and B are $n \times n$ matrices, then $\det(AB) = \det(A)\det(B)$. In particular, A is invertible if and only if $\det(A) \neq 0$ and in this case, $\det(A^{-1}) = \frac{1}{\det(A)}$.

Proof. Let $b_1, \ldots, b_n$ be the columns of B. Then

$$AB = [Ab_1 \ Ab_2 \ \cdots \ Ab_n].$$

That is, the columns of AB are $Ab_1, \ldots, Ab_n$.

Let b_j^i denote the elements of B and a_j the columns of A. Note that $Ae_j = a_j$. By linearity of the determinant as proved above we have

$$\begin{aligned}
\det(AB) = \det([Ab_1 \ Ab_2 \ \cdots \ Ab_n]) &= \det\left(\left[\sum_{j=1}^n b_1^j a_j \ Ab_2 \ \cdots \ Ab_n\right]\right) \\
&= \sum_{j=1}^n b_1^j \det([a_j \ Ab_2 \ \cdots \ Ab_n]) \\
&= \sum_{1 \leq j_1, \ldots, j_n \leq n} b_1^{j_1} b_2^{j_2} \cdots b_n^{j_n} \det([a_{j_1} \ a_{j_2} \ \cdots \ a_{j_n}]) \\
&= \left(\sum_{(j_1, \ldots, j_n) \in S_n} b_1^{j_1} b_2^{j_2} \cdots b_n^{j_n} \operatorname{sgn}(j_1, \ldots, j_n) \right) \det([a_1 \ a_2 \ \cdots \ a_n]).
\end{aligned}$$

In the above, we note that we could go from all integers, to just elements of S_n by noting that the determinant of the resulting matrix is just zero.

The conclusion follows by recognizing the determinant of B. Actually the rows and columns are swapped, but a moment's reflection will reveal that it does not matter. We could also just plug in $A = I$.

For the second part of the theorem note that if A is invertible, then $A^{-1}A = I$ and so $\det(A^{-1})\det(A) = 1$. If A is not invertible, then the columns are linearly dependent. That is suppose that

$$\sum_{j=1}^n c^j a_j = 0.$$

Without loss of generality suppose that $c^1 \neq 1$. Then take

$$B = \begin{bmatrix} c^1 & 0 & 0 & \cdots & 0 \\ c^2 & 1 & 0 & \cdots & 0 \\ c^3 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c^n & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

It is not hard to see from the definition that $\det(B) = c^1 \neq 0$. Then $\det(AB) = \det(A)\det(B) = c^1 \det(A)$. Note that the first column of AB is zero, and hence $\det(AB) = 0$. Thus $\det(A) = 0$. $\square$

Proposition: Determinant is independent of the basis. In other words, if B is invertible then,

$$\det(A) = \det(B^{-1}AB).$$

The proof is immediate. If in one basis A is the matrix representing a linear operator, then for another basis we can find a matrix B such that the matrix $B^{-1}AB$ takes us to the first basis, apply A in the first basis, and take us back to the basis we started with. Therefore, the determinant can be defined as a function on the space $L(\mathbb{R}^n)$, not just on matrices. No matter what basis we choose, the function is the same. It follows from the two propositions that

$$\det\colon L(\mathbb{R}^n) \to \mathbb{R}$$

is a well defined and continuous function.

We can now test whether a matrix is invertible.

Definition: Let $U \subset \mathbb{R}^n$ and $f\colon U \to \mathbb{R}^n$ be a differentiable mapping. Then define the *Jacobian* of f at x as

$$J_f(x) = \det\bigl(f'(x)\bigr)$$

Sometimes this is written as

$$\frac{\partial(f^1, \ldots, f^n)}{\partial(x^1, \ldots, x^n)}.$$

To the uninitiated this can be a somewhat confusing notation, but it is useful when you need to specify the exact variables and function components used.

When f is C^1, then $J_f(x)$ is a continuous function.

The Jacobian is a real valued function, and when $n = 1$ it is simply the derivative. Also note that from the chain rule it follows that:

$$J_{f\circ g}(x) = J_f\bigl(g(x)\bigr)\,J_g(x).$$

We can restate the inverse function theorem using the Jacobian. That is, $f\colon U \to \mathbb{R}^n$ is locally invertible near x if $J_f(x) \neq 0$.

For the implicit function theorem the condition is normally stated as

$$\frac{\partial(f^1, \ldots, f^n)}{\partial(y^1, \ldots, y^n)}(x_0, y_0) \neq 0.$$

It can be computed directly that the determinant tells us what happens to area/volume. Suppose that we are in $\mathbb{R}^2$. Then if A is a linear transformation, it follows by direct computation that the direct image of the unit square $A([0,1]^2)$ has area $|\det(A)|$. Note that the sign of the determinant determines "orientation". If the determinant is negative, then the two sides of the unit square will be flipped in the image. We claim without proof that this follows for arbitrary figures, not just the square.

Similarly, the Jacobian measures how much a differentiable mapping stretches things locally, and if it flips orientation. We should see more of this geometry next semester.