Index

Glossary

Linear Algebra

Definition. A *real vector space* V is a set equipped with an addition operation $+$, and a scalar ($\mathbb{R}$) multiplication operation $\cdot$ satisfying the usual identities.

Definition. A subset $W \subset V$ is a *subspace* if W is a vector space under the same operations as V.

Definition. A *linear combination* of a subset $S \subset V$ is a sum of the form

$$\sum_{v \in S} a_v \cdot v, \text{ where each } a_v \in \mathbb{R}, \text{ and only finitely many of the } \{a_v\}_{v \in S} \text{ are nonzero.}$$

Definition. The *span* of a subset $S \subset V$ is the subspace $\operatorname{span}(S)$ of linear combinations of S.

Definition. A subset $S \subset V$ is *linearly independent* if

$$\sum_{v \in S} a_v \cdot v = 0 \implies \text{each } a_v = 0.$$

Definition. A subset $S \subset V$ is a *basis* if it is linearly independent and spans V.

We assume the axiom of choice. Every vector space has a basis.

Definition. If a vector space has a finite basis of size n, we say that $\dim V = n$. Otherwise, we say that V is infinite-dimensional.

This makes sense because if $\dim V = n$, then all bases have size n.

If $W \subset V$ is a subspace, and if $\dim W = \dim V$, then $V = W$.

Definition. A map $L : V \to W$ between two vector spaces V and W is *linear* if

$$L(\alpha v_1 + \beta v_2) = \alpha L(v_1) + \beta L(v_2).$$

Definition. The *kernel* or *nullspace* of a linear map $L : V \to W$ is the subspace $\ker L \subset V$ defined by

$$\ker(L) = \{v \in V \mid L(v) = 0\}.$$

Definition. A linear map $L : V \to W$ is *injective* if $\ker L = 0$.

Definition. The *image* or *range* of a linear map $L : V \to W$ is the subspace $\operatorname{image}(L) \subset W$ defined by

$$\operatorname{image}(L) = \{L(v) \in W \mid v \in V\}.$$

Definition. A linear map $L : V \to W$ is *surjective* if $\operatorname{image}(L) = W$.

Definition. A linear map $L : V \to W$ is an *isomorphism* if there exists a linear map $L^{-1} : W \to V$ such that $L^{-1}L = \mathrm{Id}_V$ and $LL^{-1} = \mathrm{Id}_W$.

A linear map is an isomorphism iff it is both injective and surjective.

Definition. Vector spaces V and W are *isomorphic* if there exists an isomorphism from V to W.

Note that V and W are isomorphic iff $\dim V = \dim W$.

Definition. A sequence of linear maps

$$\cdots \stackrel{L_{-2}}{\to} V_{-1} \stackrel{L_{-1}}{\to} V_0 \stackrel{L_0}{\to} V_1 \stackrel{L_1}{\to} V_2 \stackrel{L_2}{\to} V_3 \stackrel{L_3}{\to} \cdots$$

is *exact* if $\ker L_i = \text{image} L_{i-1} \subset V_i$. A finite sequence of the form

$$V_0 \stackrel{L_0}{\to} V_1 \stackrel{L_1}{\to} V_2 \stackrel{L_2}{\to} \cdots \stackrel{L_{n-1}}{\to} V_n$$

is exact if this condition holds wherever it makes sense, i.e. for $i \in \{1, \ldots, n-1\}$.

Definition. Given a subspace $W \subset V_0$, a *resolution* of the subspace W is an exact sequence of the form

$$V_0 \stackrel{L_0}{\to} V_1 \stackrel{L_1}{\to} V_2 \stackrel{L_2}{\to} \cdots \stackrel{L_{n-1}}{\to} V_n \stackrel{L_n}{\to} 0$$

such that $\ker L_0 = W$.

Example. By the Poincaré lemma, the de Rham complex of a star-shaped region $U \subset \mathbb{R}^n$ is a resolution of the subspace of constant functions $\mathbb{R} \subset \Omega^0(U)$.

Theorem. *If*

$$0 \to V_1 \stackrel{L}{\to} V_2$$

is exact, then L is injective.

Proof.

$$\ker L = V_0 = 0.$$

□

Theorem. *If*

$$V_1 \stackrel{L}{\to} V_2 \to 0$$

is exact, then L is surjective.

Proof.

$$\text{image} L = \ker(V_2 \to 0) = V_2,$$

so L is surjective. □

Corollary. *If*

$$0 \to V_1 \stackrel{L}{\to} V_2 \to 0$$

is exact, then L is an isomorphism.

Proof. Since L is both injective and surjective, L is an isomorphism. □

Definition. A *short exact sequence* is an exact sequence of the form

$$0 \to V_1 \xrightarrow{L_1} V_2 \xrightarrow{L_2} V_3 \to 0.$$

Theorem. *For any short exact sequence of finite-dimensional vector spaces,* $\dim V_2 = \dim V_1 + \dim V_3$.

Proof. This proof may seem elaborate, but it follows the usual recipes of linear algebra. Suppose $\dim V_1 = n$. Then we may take a basis $\{e_i\}_{i=1}^n$ of V_1. Thus $\{L_1(e_i)\}_{i=1}^n$ span the image of L_1. Since L_1 is injective, $\{L_1(e_i)\}_{i=1}^n$ are independent in V_2. We can find vectors $\{f_j\}_{j=1}^r$ so that $\{L_1(e_i)\}_{i=1}^n \cup \{f_j\}_{j=1}^r$ becomes a basis of V_2. Thus $\dim V_2 = n + r$. Note that $\{L_1(e_i)\}_{i=1}^n \subset \text{image}(L_1) = \ker(L_2)$, so $L_2(L_1(e_i)) = 0$. Now L_2 is surjective, so the $\{L_2(f_j)\}_{j=1}^r$ must span V_3. If we can show that $\{L_2(f_j)\}_{j=1}^r$ are independent, then $\dim V_3 = r$, and the theorem follows.

To show independence, suppose $\sum a_j L_2(f_j) = 0$. Then $0 = L_2\left(\sum a_j f_j\right) \implies \sum a_j f_j \in \ker L_2 = \text{image}(L_1)$. Since $\{L_1(e_i)\}_{i=1}^n$ span the image of L_1, we have

$$\sum_{j=1}^r a_j f_j = \sum_{i=1}^n b_i L_1(e_i),$$

for some constants $\{b_i \in \mathbb{R}\}_{i=1}^n$. In particular,

$$\sum_{i=1}^n -b_i L_1(e_i) + \sum_{j=1}^r a_j f_j = 0.$$

Since $\{L_1(e_i)\}_{i=1}^n \cup \{f_j\}_{j=1}^r$ are independent in V_2, all the coefficients $\{b_i\}$ and $\{a_j\}$ must be zero. We wanted to show that all the $a_j = 0$, so we are done. □

Theorem. *If* $L : V \to W$, *then*

$$0 \to \ker L \xrightarrow{i} V \xrightarrow{L} \text{image}(L) \to 0$$

is a short exact sequence, where $i : \ker L \to V$ *is the inclusion map.*

Corollary (Rank-nullity)**.** *If* $L : V \to W$ *is a linear map of finite-dimensional vector spaces, then*

$$\dim V = \dim \ker L + \dim \text{image} L.$$

Definition. The *direct sum* of $V \oplus W$ of two vector spaces is the vector space of pairs $\{(v, w) \mid v \in V, w \in W\}$, where the operations are $(v_1, w_1) + (v_2, w_2) := (v_1 + v_2, w_1 + w_2)$ and $\alpha(v_1, w_1) = (\alpha v_1, \alpha w_1)$.

Theorem. *If* $A, B \subset V$ *are subspaces such that* $V = \text{span}(A \cup B)$, *then*

$$0 \to A \cap B \to A \oplus B \to V \to 0$$

is a short exact sequence, where $A \cap B \to A \oplus B$ *is the diagonal map* $v \mapsto (v, v)$, *and the map* $A \oplus B \to V$ *is the difference map* $(a, b) \mapsto a - b$.

Corollary.

$$\dim V = \dim A + \dim B - \dim(A \cap B).$$

Definition. If $W \subset V$ is a vector subspace, then the *quotient space* is the set of equivalence classes

$$V/W := \{[v] \mid v \in V \text{ and } [v_1] = [v_2] \iff v_1 - v_2 \in W\}.$$

The equivalence classes are thus of the form $[v] = v + W$, i.e. translates of the subspace W.

Example. Suppose $W := \{(x_1, 0, 0) \in \mathbb{R}^3 \mid x_1 \in \mathbb{R}\}$. Then $\mathbb{R}^3/W = \{[(x_1, x_2, x_3)]\}$ where we ignore the x_1-coordinate. The x_1-coordinate is still there, but it gets ignored. For example,

$$[(5, 3, 7)] = [(0, 3, 7)].$$

We could always set the x_1-coordinate to zero, but we don't have to.

Theorem. *If $W \subset V$ is a subspace, then*

$$0 \to W \xrightarrow{i} V \xrightarrow{q} V/W \to 0$$

is exact, where i is the inclusion map, and q is the quotient map $q(v) = [v]$.

Corollary.

$$\dim V/W = \dim V - \dim W.$$

Suppose that X and Y are sets, and that X has an equivalence relation $\sim$. Given an element $x \in X$, let $[x]$ denote its equivalence class. Let $X/\sim$ denote the "quotient set" of equivalence classes.

Definition. A function $f : X \to Y$ is *well-defined* on the quotient set if

$$f(x_1) = f(x_2) \text{ whenever } x_1 \sim x_2.$$

When f is well-defined on $X/\sim$, it makes sense to define a function $f : (X/\sim) \to Y$ according to the rule

$$f([x]) := [f(x)].$$

Example. Suppose V_1 and V_2 are vector spaces, and $W \subset V_1$ is a subspace. In order to define a linear map

$$V_1/W \to V_2,$$

it suffices to define a linear map

$$L : V_1 \to V_2$$

and show that L is well-defined on the quotient space. We must check the condition

$$\begin{aligned} & L(v_1) = L(v_2) \text{ whenever } v_1 \sim v_2 \\ \iff & L(v_1 - v_2) = 0 \text{ whenever } v_1 - v_2 \in W \\ \iff & L(w) = 0 \text{ whenever } w \in W \\ \iff & W \subset \ker L. \end{aligned}$$

Since W is the subspace of elements which represent zero in the quotient,

a linear map is well-defined iff everything which represents zero actually maps to zero.

Whenever this holds, it makes sense to define $L : (V_1/W) \to V_2$ by

$$L([v]) := [L(v)].$$

Chain complexes

Definition. A sequence of linear maps

$$\cdots \overset{d^{p-2}}{\to} C^{p-1} \overset{d^{p-1}}{\to} C^p \overset{d^p}{\to} C^{p+1} \overset{d^{p+1}}{\to} C^{p+2} \overset{d^{p+2}}{\to} C^{p+3} \overset{d^{p+3}}{\to} \cdots$$

is called a *cochain complex* if image $d^{p-1} \subset \ker d^p \subset C^p$. Equivalently, $d^p \circ d^{p-1} = 0$ for each p. The vector spaces C^p are called *cochains*, and the linear maps d^p are called *differentials* or *coboundary maps*.

Note that every exact sequence is a cochain complex, but not conversely. The difference is that for an exact sequence, we demand that image $d^{p-1} = \ker d^p$.

Example. The de Rham complex $\Omega^\bullet(U)$ of an open subset $U \subset \mathbb{R}^n$ is a cochain complex with $C^p = \Omega^p(U)$ and d^p the exterior derivative d.

Definition. The *cocycles* of a cochain complex C are the subspaces $Z^p = Z^p(C) := \ker d^p \subset C^p$.

(The letter Z comes from the German word “zyklus” for cycle.)

Example. In the de Rham complex, cocycles are called *closed forms*.

$$Z^p(U) := \{\omega \in \Omega^p(U) \mid d\omega = 0\} = \ker(d : \Omega^p(U) \to \Omega^{p+1}(U)) \subset \Omega^p(U).$$

Definition. The *coboundaries* of a cochain complex C are the subspaces $B^p = B^p(C) :=$ image $d^{p-1} \subset Z^p(C) \subset C^p$.

Example. In the de Rham complex, coboundaries are called *exact forms*.

$$B^p(U) := \left\{d\eta \mid \eta \in \Omega^{p-1}(U)\right\} = \text{image}\left(d : \Omega^{p-1}(U) \to \Omega^p(U)\right).$$

Definition. The *cohomology* of a cochain complex C is the collection of quotient spaces $H^p = H^p(C) := Z^p(C)/B^p(C) = \ker d^p/\text{image}\, d^{p-1}$.

Example. If $U \subset \mathbb{R}^n$ is open, then the cohomology of the de Rham complex is

$$H^p(U) := H^p(\Omega^\bullet(U)) = \frac{\{\omega \in \Omega^p(U) \mid d\omega = 0\}}{\{d\eta \mid \eta \in \Omega^{p-1}(U)\}}.$$

Any vector $[\omega] \in H^p(U)$ is represented by some $\omega \in \Omega^p(U)$, and $[\omega_1] = [\omega_2]$ iff $\omega_1 = \omega_2 + d\eta$ for some η.

Definition. Given two chain complexes A and B, a chain map $f : A \to B$ is a collection of maps $\{f^p : A^p \to B^p\}$ such that $d^p \circ f^p = f^{p+1} \circ d^p$.

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{d^{p-1}} & A^p & \xrightarrow{d^p} & A^{p+1} & \xrightarrow{d^{p+1}} & A^{p+2} & \xrightarrow{d^{p+2}} & \cdots \\
 & & \downarrow{\scriptstyle f^p} & & \downarrow{\scriptstyle f^{p+1}} & & \downarrow{\scriptstyle f^{p+2}} & & \\
\cdots & \xrightarrow{d^{p-1}} & B^p & \xrightarrow{d^p} & B^{p+1} & \xrightarrow{d^{p+1}} & B^{p+2} & \xrightarrow{d^{p+2}} & \cdots
\end{array}$$

Example. Suppose $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are open subsets, and $\phi : U \to V$ is C^∞. Then ϕ^* is a chain map of the de Rham complexes:

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{d^{p-1}} & \Omega^p(V) & \xrightarrow{d} & \Omega^{p+1}(V) & \xrightarrow{d} & \Omega^{p+2}(V) & \xrightarrow{d} & \cdots \\
 & & \downarrow{\scriptstyle \phi^*} & & \downarrow{\scriptstyle \phi^*} & & \downarrow{\scriptstyle \phi^*} & & \\
\cdots & \xrightarrow{d^{p-1}} & \Omega^p(U) & \xrightarrow{d} & \Omega^{p+1}(U) & \xrightarrow{d} & \Omega^{p+2}(U) & \xrightarrow{d} & \cdots
\end{array}$$

Theorem. *If $f : A \to B$ is a cochain map, then f determines a well-defined map (usually also denoted by f) on cohomology.*

$$\begin{array}{ccccc}
\cdots & H^{p-1}(A) & H^p(A) & H^{p+1}(A) & \cdots \\
 & {\scriptstyle f}\downarrow & {\scriptstyle f}\downarrow & {\scriptstyle f}\downarrow & \\
\cdots & H^{p-1}(B) & H^p(B) & H^{p+1}(B) & \cdots
\end{array}$$

Example. Suppose $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are open subsets, and $\phi : U \to V$ is C^∞. Then ϕ^* gives maps

$$\begin{array}{ccccc}
\cdots & H^p(V) & H^{p+1}(V) & H^{p+2}(V) & \cdots \\
 & \downarrow{\scriptstyle \phi^*} & \downarrow{\scriptstyle \phi^*} & \downarrow{\scriptstyle \phi^*} & \\
\cdots & H^p(U) & H^{p+1}(U) & H^{p+2}(U) & \cdots
\end{array}$$

Definition. Given two cochain maps $f, g : A \to B$, a *cochain homotopy* from f to g is a collection of maps $S = \{S^p : A^p \to B^{p-1}\}$ such that

$$dS + Sd = f - g,$$

i.e. $d^{p-1}S^p + S^{p+1}d^p = f^p - g^p$ for each p. Two cochain maps f and g are said to be *cochain homotopic* or *homotopic* if there exists

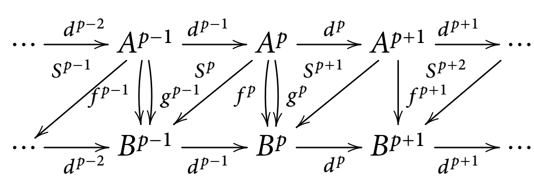

Definition. A smooth homotopy between two smooth maps $\phi_0, \phi_1 : U \to V$ is a smooth map $\phi : [0,1] \times U \to V$ such that $\phi|_{t=0} = \phi_0$ and $\phi|_{t=1} = \phi_1$.

Example. Suppose $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are open subsets, and suppose $\phi : [0,1] \times U \to V$ is any smooth map, so that ϕ_0 is smoothly homotopic to ϕ_1. The homotopy operator S_ϕ gives a cochain homotopy from ϕ_1^* to ϕ_0^*.

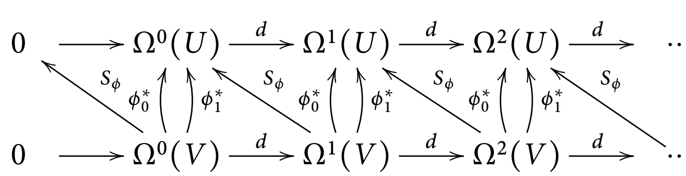

Theorem. *If two cochain maps are cochain homotopic, then they determine the same map on cohomology.*

Proof. Suppose f and g are cochain homotopic. Then

$$f([x]) - g([x]) = [(f-g)(x)] = \left[dSx + S\cancel{dx}\right] = \cancel{\left[dSx\right]} = 0.$$

□

Example. Suppose $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are open subsets, and $\phi : [0,1] \times U \to V$ is a smooth homotopy, so that ϕ_0 and ϕ_1 are smoothly homotopic. Then $\phi_0^* = \phi_1^*$ on cohomology:

$$\begin{array}{ccccc}
\cdots & H^p(V) & H^{p+1}(V) & H^{p+2}(V) & \cdots \\
 & \downarrow{\scriptstyle \phi_0^*=\phi_1^*} & \downarrow{\scriptstyle \phi_0^*=\phi_1^*} & \downarrow{\scriptstyle \phi_0^*=\phi_1^*} & \\
\cdots & H^p(U) & H^{p+1}(U) & H^{p+2}(U) & \cdots
\end{array}$$

Dual spaces

Definition. Given a vector space V, the *dual space* is

$$V^* := \{\alpha : V \to \mathbb{R} \mid \alpha \text{ is linear}\}.$$

Suppose $\dim V = n$, so we have a finite basis $\{e_i\}_{i=1}^n$.

Definition. The *dual basis* $\{\varepsilon_i\}_{i=1}^n$ to a basis $\{e_i\}_{i=1}^n$ is characterized by the properties that the $\varepsilon_i : V \to \mathbb{R}$ are linear and satisfy

$$\varepsilon_i(e_j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

Any vector $v \in V$ can be written as $v = \sum_{i=1}^n v_i e_i$. From the above definition, we see that the ε_i are the *component functions* $\varepsilon_i(v) = v_i$. Thus for any basis and its dual, we have the general formula

$$v = \sum_{i=1}^n \varepsilon_i(v) e_i.$$

In physics, this is often called *resolution of the identity*.

Theorem. *The ε_i form a basis for the dual space.*

Proof. We need to show that the ε_i are linearly independent and span. For independence, suppose

$$\alpha := \sum_{i=1}^n a_i \varepsilon_i, \text{ and } \alpha = 0.$$

Then for each i we have both $\alpha(e_i) = a_i$ from the definition of the dual basis, and $\alpha(e_i) = 0$ from the assumption $\alpha = 0$. Therefore, each $a_i = 0$, so there is no nontrivial linear dependence among the ε_i.

To show that the ε_i span, consider some $\beta \in V^*$. Define $b_i := \beta(e_i)$ so that each $b_i \in \mathbb{R}$. Thus

$$\tilde{\beta} := \sum b_i \varepsilon_i \in V^*.$$

I claim that $\beta = \tilde{\beta}$. To show that two functions are equal, it suffices to prove that they agree at each point. A general point $v \in V$ can be written as $v = \sum v_i e_i$. Now we compare

$$\beta(v) = \beta\left(\sum v_i e_i\right) = \sum v_i \beta(e_i) = \sum v_i b_i,$$
$$\tilde{\beta}(v) = \left(\sum b_i \varepsilon_i\right)(v) = \sum b_i \varepsilon_i(v) = \sum b_i v_i.$$

Thus $\beta(v) = \tilde{\beta}(v)$ for all $v \in V$, so $\beta = \tilde{\beta}$. Since $\tilde{\beta}$ is a linear combination of $\{\varepsilon_i\}$, we conclude that $\{\varepsilon_i\}$ spans V^*. Since the $\{\varepsilon_i\}$ are linearly independent and span, we conclude that $\{\varepsilon_i\}$ is a basis. □

Calculus

Definition. An *open set* $U \subset \mathbb{R}^n$ is a set such that for every point $x \in U$, there is some positive radius r (depending on x) such that the open ball

$$B_r(\vec{x}) := \{y \in \mathbb{R}^n \mid |x - y| < r\}$$

is contained in U. Intuitively, every point in U is an interior point.

Define $V := \mathbb{R}^n$. Consider the standard basis $\{e_i\}_{i=1}^n$ of $\mathbb{R}^n$, so that $e_1 = (1, 0, \ldots, 0)$, $e_2 = (0, 1, 0, \ldots, 0)$, etc. Then any vector $v = (v_1, \ldots v_n) \in V$ can be written as

$$v = \sum_{i=1}^n v_i e_i.$$

Definition. Given a function $f : U \to \mathbb{R}$ and a vector $v \in V$, the *directional derivative* $\partial_v f$ is the function defined by

$$(\partial_v f)(x) := \lim_{h \to 0} \frac{f(x + hv) - f(x)}{h},$$

on the subdomain where this limit exists.

Definition. The *partial derivatives* are

$$\frac{\partial f}{\partial x_i} := \partial_{e_i} f.$$

Definition. For a nonnegative integer k, a function $f : U \to \mathbb{R}$ is of *class* C^k if all iterated partial derivatives of f up to order k both exist over all of U, and are continuous. We write $f \in C^k(U, \mathbb{R})$, or $f \in C^k(U)$.

Definition. A function $f : U \to \mathbb{R}$ is of *class* C^∞ if $f \in C^k$ for all k.

Definition. The *coordinate functions* $\{x_i\}_{i=1}^n$ are functions $x_i \in C^\infty(U, \mathbb{R})$ defined so that $x_i(v) = \varepsilon_i(v) = v_i$ for any $v \in U$.

Even though $x_i = \varepsilon_i$ as functions, we regard them as living in different vector spaces. Specifically, $x_i \in C^\infty(U, \mathbb{R})$, while $\varepsilon_i \in V^*$.

Suppose W is a finite-dimensional vector space with basis $\{w_j\}_{j=1}^m$. Given any function $\phi : U \to W$, we can express ϕ in terms of component functions

$$\phi = \sum \phi_j w_j,$$

where $\phi_j : U \to \mathbb{R}$. In this particular basis, $\phi = (\phi_1, \dots \phi_m)$.

Definition. $C^k(U, W)$ and $C^\infty(U, W)$ are the spaces of functions $\phi : U \to W$ such that each component function ϕ_j belongs to $C^k(U, \mathbb{R})$ or $C^\infty(U, \mathbb{R})$ respectively.

Definition. If $U_1 \subset V_1$ and $U_2 \subset V_2$ are open subsets, then $C^k(U_1, U_2)$ and $C^\infty(U_1, U_2)$ are the respective *subsets* of functions ϕ in $C^k(U_1, V_2)$ and $C^\infty(U_1, V_2)$ such that $\phi(U_1) := \{\phi(x) \in V_2 \mid x \in U_1\} \subset U_2$.

Warning $C^k(U_1, U_2)$ and $C^\infty(U_1, U_2)$ are usually not vector spaces!

The building blocks of multivariable functions $\phi : U_1 \to U_2$ are component functions $\phi_j : U_1 \to \mathbb{R}$, so these are the most important functions to understand.

Definition. For an open subset $U \subset V \cong \mathbb{R}^n$, then the *total derivative* of a function $f : U \to \mathbb{R}$ is

$$Df := \frac{\partial f}{\partial x_1}\varepsilon_1 + \cdots + \frac{\partial f}{\partial x_n}\varepsilon_n.$$

If $f \in C^k(U, \mathbb{R})$ for $k > 0$, then $Df \in C^{k-1}(U, V^*)$.

Theorem. *If $f \in C^1(U, \mathbb{R})$, then*

$$\partial_v f = Df(v) = \sum_{i=1}^n v_i \frac{\partial f}{\partial x_i}.$$

This theorem tells us that for C^1 functions, the partial derivatives determine the directional derivatives.

Theorem. *If $f \in C^2(U, \mathbb{R})$, then*

$$\partial_v \partial_w f = \partial_w \partial_v f$$

for all $v, w \in V$. In particular, the partial derivatives are symmetric:

$$\frac{\partial}{\partial x_i}\frac{\partial}{\partial x_j} f = \frac{\partial}{\partial x_j}\frac{\partial}{\partial x_i} f.$$

The proof of this theorem is surprisingly subtle. The standard proof uses Green's theorem (see Spivak). Alternatively, it is possible to prove this at any point by using a second-order Taylor polynomial with remainder. (The theorem is obvious on polynomials, and one shows that the remainder does not contribute.)

Theorem. *Suppose U and V are open subsets of $\mathbb{R}^n$, and $\phi \in C^1(U, V)$ has a C^0 inverse. Suppose $f \in C^0(V)$ has compact support. Then*

$$\int_V f(v)\, dx_1 \cdots dx_n = \int_U (\phi^* f)(u)\ |\det(D\phi)(u)|\ dx_1 \cdots dx_n < \infty.$$

Topology

Definition. A *topology* on a set X is a collection of subsets $\mathcal{T}$ called "open sets" which satisfy

- The empty set $\varnothing$ and the whole set X are open,
- The union of an arbitrary collection of open sets is open,
- The intersection of a finite collection of open sets is open.

Definition. A pair $(X, \mathcal{T})$ consisting of a set X and a topology $\mathcal{T}$ on X is called a *topological space*. When the topology $\mathcal{T}$ is clear from context, we refer to X as the topological space.

Definition. A subset C of a topological space X is *closed* if the complement $X - C$ is open.

Example. $\mathbb{R}^n$ with the standard definition of "open set" is a topological space. This is called the *standard topology* on $\mathbb{R}^n$. The closed sets of $\mathbb{R}^n$ agree with the standard definition.

Example. The *trivial topology* on a set X is the topology $\mathcal{T} = \{\varnothing, X\}$.

Example. The *discrete topology* on a set X is $\mathcal{T} = \mathcal{P}(X) := \{\text{all subsets of } X\}$.

Definition. Given a topological space $(X, \mathcal{T})$, the *relative topology, subspace topology,* or *induced topology* of a subspace $S \subset X$ is the topology on S defined so that the open sets are

$$\{S \cap U \mid U \in \mathcal{T}\}.$$

Example. Consider the unit two-sphere $S^2 \subset \mathbb{R}^3$. The standard topology on $\mathbb{R}^3$ induces a subspace topology on S^2. In that induced topology, a subset $U \subset S^2$ is open iff its stereographic projections are open.

Definition. The *interior* $\mathring{S}$ of a subset S in a topological space $(X, \mathcal{T})$ is the union of all open sets contained in S, i.e.

$$\mathring{S} := \bigcup \{U \in \mathcal{T} \mid U \subset S\}.$$

Similarly, the *closure* $\overline{S}$ is the intersection of all closed sets containing S.

Definition. An *open cover* of a subset $S \subset X$ of a topological space $(X, \mathcal{T})$ is a collection of open subsets $\mathcal{U} \subset \mathcal{T}$ such that S is contained in the union, i.e.

$$S \subset \bigcup \mathcal{U}.$$

Definition. A subset $K \subset X$ of a topological space is *compact* if every open cover of K has a finite subcover.

Example. A subset $K \subset \mathbb{R}^n$ is compact iff it is closed and bounded (Heine-Borel theorem).

For any positive δ, let $\mathcal{U}_\delta$ denote the set of all balls in $\mathbb{R}^n$ of radius δ. Note that $\mathcal{U}_\delta$ covers all of $\mathbb{R}^n$. If $K \subset \mathbb{R}^n$ is closed and bounded, then $\mathcal{U}_\delta$ is certainly a cover of K. By the compactness of K, we can find finitely many balls of radius δ which cover K.

Definition. A topological space X is *Hausdorff* if any two distinct points $x_1, x_2 \in X$ are contained in disjoint open sets.

Example. The trivial topology on any set X with two or more elements is not Hausdorff since the only open set containing x_1 is X itself, which also contains x_2.

Definition. A *basis* for a topological space $(X, \mathcal{T})$ is a collection $B \subset \mathcal{T}$ of open sets such that any open subset $U \subset X$ can be written as a union of elements of U.

Example. The set of open balls is a basis for the standard topology on $\mathbb{R}^n$.

Definition. A topological space is *second-countable* if it has a countable basis.

Example. The set of open balls with rational centers and rational radii form a countable basis for the standard topology on $\mathbb{R}^n$. Therefore, the standard topology on $\mathbb{R}^n$ is second-countable.

Definition. A map $f : X \to Y$ between two topological spaces is *continuous* if the inverse image of any open set is open.

Theorem. *Using the standard topology, any map $f : \mathbb{R}^m \to \mathbb{R}^n$ is continuous in the sense of topological spaces iff it is continuous in the sense of δ and ε.*

Proof. First suppose that f is δ-ε-continuous. We must show that the inverse image of any open set $V \subset \mathbb{R}^n$ is open. Equivalently, for each $x_0 \in f^{-1}(V)$ we must find an open ball $B_\delta(x_0)$ of radius δ and centered at x_0 such that $B_\delta(x_0) \subset f^{-1}(V)$. Consider $f(x_0) \in V$. Since V is open, we can find a ball $B_\varepsilon(f(x_0))$ with radius ε and center $f(x_0)$ which is contained in V. By δ-ε-continuity, there exists $\delta > 0$ such that

$$|x - x_0| < \delta \implies |f(x) - f(x_0)| < \varepsilon,$$

i.e. $f(B_\delta(x_0)) \subset B_\varepsilon(f(x_0))$. It follows that

$$B_\delta(x_0) \subset f^{-1}(B_\varepsilon(f(x_0)) \subset f^{-1}(V).$$

We have found our required open ball $B_\delta(x_0)$ around the general point $x_0 \in f^{-1}(V)$, so we know that the inverse image of an open subset is open.

It remains to show the converse: if f is continuous in the sense of topological spaces, then f is δ-ε-continuous. For any $x_0 \in \mathbb{R}^m$ and $\varepsilon > 0$. We wish to produce $\delta > 0$ such that $|x - x_0| < \delta \implies |f(x) - f(x_0)| < \varepsilon$. Equivalently,

$$f(B_\delta(x_0)) \subset B_\varepsilon(f(x_0)).$$

for some $\delta > 0$. It suffices to show that

$$B_\delta(x_0) \subset f^{-1}(B_\varepsilon(f(x_0))$$

for some δ, since $f(f^{-1}(V)) = V$ for any subset $V \subset \mathbb{R}^n$. (Compare with $f^{-1}(f(U)) \supset U$, where equality is not generally true.)

Now we are in position to use the definition of continuity for topological spaces. Since $B_\varepsilon(f(x_0))$ is open, by the definition of continuity for topological spaces, the inverse image $f^{-1}(B_\varepsilon(f(x_0)))$ is also open. Clearly $x_0 \in f^{-1}(f(x_0)) \subset f^{-1}(B_\varepsilon(f(x_0)))$. Therefore, we can find some open ball of positive radius around x_0 which is contained in $f^{-1}(B_\varepsilon(f(x_0)))$. The radius of this ball is our desired δ. □

Definition. Let X be a topological space, and suppose $f : X \to \mathbb{R}$. The *support* of f, denoted $\operatorname{supp}(f)$, is the closure of $\{x \in X \mid f(x) \neq 0\}$.

Example. Suppose ψ is a cutoff function so that $\psi(x) = 0$ for $x \leq 0$, and $\psi(x) > 0$ for $x > 0$. Then $\operatorname{supp}(\psi) = \overline{(0, \infty)} = [0, \infty)$.

Definition. A function $f : X \to \mathbb{R}$ has *compact support* if $\operatorname{supp}(f)$ is compact.

Warning Suppose $U \subset \mathbb{R}^n$ is open, and $f \in C^0(U)$. The support of f is computed using the *subspace topology* of U. For f to have compact support, it does not suffice to have $\{x \in X \mid f(x) \neq 0\}$ be a bounded subset of $\mathbb{R}^n$.

Theorem. *If $U \subset \mathbb{R}^n$ and $f \in C^0(U)$, then f has compact support iff the closure in $\mathbb{R}^n$ of $\{x \in \mathbb{R}^n \mid x \in U \text{ and } f(x) \neq 0\}$ is both bounded and contained in U.*