Crash course: constructions in point-set topology and linear algebra

Keeley Hoek*

May 25, 2022

The purpose of these notes is to provide an introduction to fundamental constructions in topology and linear algebra, in order to prepare us to combine them as we perform basic constructions on vector bundles. As we go, we gently emphasize the so-called "universal properties" which these constructions satisfy.

On the topology side, primarily as a source of examples we assume knowledge of the basics of metric spaces (but not general topological spaces).[1] On the linear algebra side, we assume knowledge of free vector spaces and vector space quotients (neither of these things are particularly scary if they are new to you—don't worry!).

In general, the content spans several different levels and these notes should be a useful reference throughout the class (as opposed to being intended for you to completely memorize). Hopefully there is something here for everyone; for the most part we include the fundamentals and basic definitions for people just getting started, but even those familiar with much of the content might be interested in the slightly more abstract formulation given here than one might see in an introductory class. There are also completely optional (clearly marked) advanced pieces which appear towards the end of each section, and which won't have anything to do with our tutorial proper.

There are lots of exercises with varying difficulty, and many demonstrate a deep or important fact which is good to know on its own. Thus, it will be useful for you to try to understand them all even if you will not try to solve them all. As a general rule the exercises get harder and less important as each section goes on. The harder exercises which are still important have been split into multiple sub-parts to make them more manageable.

Examples appear very sparsely in order to keep the size under control; these notes are no substitute for a real book on topology or linear algebra, but hopefully provide a useful starting point. Feel free to ask for references or for more examples to be added!

---

*Please send questions and corrections to `khoek@math.harvard.edu`.

[1] Note that we assume significantly more prerequisite knowledge in the (totally optional) Subsection 1.5.

Contents

Notable paragraphs are highlighted as described below.

Definitions to know/be familiar with are highlighted in green.

Exercises appear in blue.

Material which you do not need for this class (and may require additional prerequisites to even understand, so don't fret) appears in red. You still might find them interesting!

1 Topology

1.1 Fundamentals

Definition 1.1. A *topology* on a set X is a collection of subsets $\mathcal{T}$ which

- is closed under arbitrary unions (including empty ones, i.e. contains $\emptyset$), and
- is closed under finite intersections (including empty ones, i.e. contains X).

A set X equipped with a topology $\mathcal{T}$ on X is called a *topological space* (or often just a *space*).

The subsets $U \in \mathcal{T}$ of X are called *open* (with respect to the topology $\mathcal{T}$). Likewise, a subset C of X is *closed* if its complement is open.

When no confusion about the topology arises, we will refer to a topological space $(X, \mathcal{T})$ just by the name X of its underlying set. Elements $x \in X$ of a topological space are called *points*.

Exercise 1.2. Restate the definition of a topology in terms of closed sets.

Exercise 1.3. Show that there is a unique topological space $\bullet$ with one point.

Example 1.4. The space $\mathbb{R}^n$ becomes a topological space by declaring that a subset $U \subseteq \mathbb{R}^n$ is open exactly when for each $x \in U$ there is some $\varepsilon > 0$ such that the ε-ball about x is wholly contained in U.

Exercise 1.5. Verify that this definition of the open subsets of $\mathbb{R}^n$ actually specifies a topology.

Example 1.6. If (X, d) is a metric space then X canonically becomes a topological space using the same definition as for $\mathbb{R}^n$ (the same verification works).

Definition 1.7. A function $f : X \to Y$ between topological spaces is *continuous* if the preimage $f^{-1}(U)$ of every open subset $U \subseteq X$ is open. A function $f : X \to Y$ between topological spaces is a *homeomorphism* if f is bijective, continuous, and has a continuous inverse.[a]

[a] Danger! This is not the same as requiring that f is merely a bijective continuous function.

We will often simply call a continuous function $f : X \to Y$ a *map*. If there exists a homeomorphism $f : X \to Y$ between topological spaces we say that they are *homeomorphic*. Homeomorphic topological spaces are "the same": a homeomorphism $f : X \to Y$ induces both a correspondence between the points of X and Y, and a compatible correspondence between the open sets of X and Y.

Exercise 1.8. Verify that a function $f : (X, d) \to (Y, d')$ between metric spaces is a continuous map $X \to Y$ of topological spaces if and only if f satisfies the usual ε-δ definition of continuity.

Definition 1.9. If $A \subseteq X$ is a subset of a topological space then its *closure* $\overline{A}$ is the intersection of all closed sets which contain A.

Exercise 1.10. If $A \subseteq X$ is a subset of a topological space show that

1. the set $\overline{A}$ is closed,
2. the set A is closed if and only if $A = \overline{A}$, and
3. if $B \subseteq A$ then $\overline{B} \subseteq \overline{A}$.

Exercise 1.11. Give an example of a chain of subsets $B \subseteq A \subseteq \mathbb{R}$ such that we have $\overline{B} \nsubseteq A$.

Exercise 1.12. Show that if (X, d) is a metric space, $A \subseteq X$ is a subset, and (x_n) is a sequence in A converging to some $x \in X$, then $x \in \overline{A}$.

1.2 Subspaces, quotients, and (co)products

Next we'll discuss some basic constructions one can perform to make new topological spaces from old ones.

Definition 1.13. Let $A \subseteq X$ be a subset of a topological space X with topology $\mathcal{T}$. If $\iota : A \hookrightarrow X$ denotes the inclusion, then

$$\mathcal{S} = \{\iota^{-1}(U) : U \in \mathcal{T}\}$$

is a topology on A called the *subspace topology*.

When we equip a subset $A \subseteq X$ of a topological space with the subspace topology, we call A a *subspace* of X.

Exercise 1.14. Let $A \subseteq X$ be a subset of a topological space X.

1. Show that the subspace topology is actually a topology.
2. Show that when A is given the subspace topology then $U \subseteq A$ is open if and only if there is an open subset $V \subseteq X$ such that $U = A \cap V$.

Proposition 1.15 (Universal property of subspaces)**.** *The subspace topology on $A \subseteq X$ is the unique topology such that for every topological space Y a function $f : Y \to A$ is continuous if and only if $\iota \circ f : Y \to A \hookrightarrow X$ is continuous.*

Exercise 1.16. Prove Proposition 1.15. (The proof is very short!)

Exercise 1.17. Show that the arbitrary intersection of topologies on a set A is still a topology on A. Conclude by Proposition 1.15 that one can define the subspace topology on $A \subseteq X$ as the intersection of all topologies for which a function $f : Y \to A$ is continuous whenever $\iota \circ f : Y \to A \hookrightarrow X$.

In fact, since arbitrary intersections of topologies are still topologies, it always makes sense to speak of the "coarsest topology" on X for which certain functions $Y \to X$ are continuous.

Example 1.18. The unit interval $I := [0, 1]$ is a subspace of $\mathbb{R}$. The Mandelbrot set is a subspace of the complex plane $\mathbb{C}$.

Subspaces are obtained by deleting points from our original topological space X, but sometimes we want to identify multiple points of X with each other instead. This is called taking a quotient.

Definition 1.19. Let $\sim$ be an equivalence relation on a topological space X with topology $\mathcal{T}$ and let $X/\sim$ be the quotient of sets. If $\pi : X \to X/\sim$ is the projection then

$$\mathcal{S} = \{U \subseteq X/\sim\, : \pi^{-1}(U) \in \mathcal{T}\}$$

is a topology on $X/\sim$ called the *quotient topology*.

Exercise 1.20. Let $\sim$ be an equivalence relation on a topological space X. Show that the quotient topology on $X/\sim$ is actually a topology.

Proposition 1.21 (Universal property of quotients)**.** *The quotient topology on $X/\sim$ is the unique topology such that for every topological space Y a function $f : X/\sim\, \to Y$ is continuous if and only if $f \circ \pi : X \twoheadrightarrow X/\sim\, \to Y$ is continuous.*

Exercise 1.22. Prove Proposition 1.21. (The proof is very short!)

We again conclude (as in Exercise 1.17) that the quotient topology on $X/\sim$ is the intersection of all topologies for which $f : X/\sim \to Y$ is continuous whenever $f \circ \pi : X \twoheadrightarrow X/\sim \to Y$ is continuous.

Example 1.23. We can define an equivalence relation $\sim$ on the interval $I = [0, 1]$ which identifies 0 and 1 and leaves all other points related only to themselves: intuitively we have "glued the point $0 \in I$ to the point $1 \in I$". The result is the circle $S^1 := I/\sim$.

Exercise 1.24. Consider the equivalence relation $\sim'$ on $\mathbb{R}$ for which $x \sim' y$ exactly when $x - y \in \mathbb{Z}$.

1. Show that the quotient $\mathbb{R}/\sim'$ gives another construction of the circle by writing down an explicit homeomorphism $\mathbb{R}/\sim' \to I/\sim$.
2. Show that S^1 is also homeomorphic to the subspace of the complex plane $\mathbb{C}$ consisting of points of modulus 1.

Exercise 1.25. Show that $\mathbb{R}$ is homeomorphic to the subspace of S^1 with the point $[0]$ deleted.[a]

[a] Here the notation $[0]$ denotes the equivalence class of 0 in the quotient $I/\sim$.

Example 1.26. There is an equivalence relation on $\mathbb{R}^n \setminus \{0\}$ defined by specifying that $x \sim'' y$ exactly when x and y are scalar multiples (that is, if x and y lie on the same line). The quotient $(\mathbb{R}^n \setminus \{0\})/\sim''$ is called *real projective n-space* and is denoted $\mathbb{R}P^n$.

Exercise 1.27. Let D^n be the unit disk in $\mathbb{R}^n$, and let S^{n-1} the the unit sphere in $\mathbb{R}^n$ (the boundary of D^n). Define an equivalence relation on D^n which identifies all points of S^{n-1} and leaves all other points unrelated. Show that the quotient of D^n by this equivalence relation is homeomorphic to the unit sphere in $\mathbb{R}^{n+1}$, i.e. S^n.

Just as we can take products $X \times Y$ of sets, we can also take products of topological spaces. If the explicit description of the topology on $X \times Y$ below is a bit mysterious, don't worry: take the definition for granted and then do Exercise 1.31 below.

Definition 1.28. Let X and Y be topological spaces with respective topologies $\mathcal{T}_X$ and $\mathcal{T}_Y$. If $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ are the projections then

$$\mathcal{S} = \left\{ \bigcup_{\alpha} \left(\pi_X^{-1}(U_\alpha) \cap \pi_Y^{-1}(V_\alpha) \right) : U_\alpha \in \mathcal{T}_X, V_\alpha \in \mathcal{T}_Y \right\}$$

is a topology on $X \times Y$ called the *product topology*.

Exercise 1.29. Show that the product topology on $X \times Y$ is actually a topology.

Proposition 1.30 (Universal property of products)**.** *The product topology on $X \times Y$ is the unique topology such that for every topological space Z a function $f : Z \to X \times Y$ is continuous if and only if $\pi_X \circ f : Z \to X \times Y \twoheadrightarrow X$ and $\pi_Y \circ f : Z \to X \times Y \twoheadrightarrow Y$ are both continuous.*

Exercise 1.31. Prove Proposition 1.30. (The proof is not as short as the previous examples, but is still not so long.)

Exercise 1.32. Explain how generalize Definition 1.28 from a pair of topological spaces X and Y to a possibly-infinite family $\{X_\alpha\}$.

Example 1.33. The space $\mathbb{R}^2$ is homeomorphic to the product $\mathbb{R} \times \mathbb{R}$, and likewise $\mathbb{R}^n$ is homeomorphic to the n-fold product of $\mathbb{R}$ with itself. (Given our definition of the product topology, this is a non-trivial fact.)

Example 1.34. The 2-torus is the product $\mathbb{T}^2 := S^1 \times S^1$ (the circle S^1 was defined in Example 1.23).

Example 1.35. If $f : X \to X$ is a homeomorphism then the *mapping torus* M_f is the quotient of the product $X \times I$ by the equivalence relation which identifies $(x, 0) \in X \times I$ with $(f(x), 1) \in X \times I$. Mapping tori play an important role in the theory of 3-manifolds.

Formally dual to products are *coproducts*: there is a canonical topology on the disjoint union $X \sqcup Y$ of two (or indeed infinitely many) topological spaces. As you likely expect by now, the topology is obtained by formally "reversing the arrows" in the universal property of products.

Proposition 1.36 (Universal property of coproducts)**.** *The coproduct topology on $X \sqcup Y$ is the unique topology such that for every topological space Z a function $f : X \sqcup Y \to Z$ is continuous if and only if $f \circ \iota_X : X \hookrightarrow X \sqcup Y \to Z$ and $f \circ \iota_Y : Y \hookrightarrow X \sqcup Y \to Z$ are both continuous.*

Exercise 1.37. Give an explicit definition of the topology on $X \sqcup Y$ so that Proposition 1.36 is true. Explain how to generalize your definition to infinite coproducts.

Exercise 1.38. Above we described (co)products of pairs of spaces, which are more accurately called *binary (co)products*. State the natural generalizations to finite and infinite sequences of spaces, and thereby define *finite* and *infinite (co)products*.

Example 1.39. A common way to build a topological space out of simple pieces is to take the coproduct of all of the pieces, and then take a quotient of the resulting space. For instance, if $f : X \to Y$ is a continuous map the *mapping cylinder* M_f is the quotient of $(X \times I) \sqcup Y$ by the equivalence relation which identifies $(x, 1) \in X \times I$ with $f(x) \in Y$ (and leaves all other points unrelated). This is a generalization of Example 1.35.

Forming mapping cylinders is a fundamental operation in homotopy theory, because the inclusion $X \hookrightarrow M_f$ defined by $x \mapsto [(x, 0)]$ is always a *cofibration*.

A closely related construction is the *mapping cone* C_f, which is the further quotient of M_f which identifies $(0, x)$ and $(0, x')$ for all $x, x' \in X$.

1.3 Compactness and paracompactness

Definition 1.40. Let X be a topological space.

- A *cover* $\mathcal{U}$ of X is a collection of subsets of X such that $\bigcup \mathcal{U} = X$.[a]
- A *subcover* $\mathcal{V}$ of a cover $\mathcal{U}$ is a subset $\mathcal{V} \subseteq \mathcal{U}$ which is still a cover.
- A *refinement* $\mathcal{V}$ of a cover $\mathcal{U}$ is another cover with the property that each $V \in \mathcal{V}$ is wholly contained in some $U \in \mathcal{U}$.

[a]That is, the union of all of the sets in $\mathcal{U}$ is X.

Note that subcovers are examples of refinements. We will almost always speak of *open covers* of X, which are covers of X consisting exclusively of open sets.

Definition 1.41. A topological space X is *compact* if every open cover of X has a finite subcover.

Example 1.42. Spaces with finitely many points are always compact. On the other hand, the family of intervals $\{(-\varepsilon, \varepsilon) : \varepsilon \in (0, 1)\}$ in $\mathbb{R}$ is an open cover of the interval $(-1, 1)$ with no finite subcover (why?).

The following theorem shows that, for metric spaces, our general notion of compactness is familiar.

Theorem 1.43. *Let (X, d) be a metric space. Then X is compact if and only if X is sequentially compact*[2].

We prove Theorem 1.43 in Subsection 1.5 below, but doing this requires introducing considerably more point-set topology machinery than we will need in this class (we won't care at all about metric spaces in particular), so all of it is optional and provided for the interested reader.

Before moving on we note a few exercises and additional facts about compact spaces which, in order to prove, also require more point-set topology than you need for this class.

[2]Recall that a metric space is sequentially compact if every sequence has a convergent subsequence.

Exercise 1.44. Show that the product $X \times Y$ of compact topological spaces is compact.

In fact, we have the following generalization of Exercise 1.44.

Theorem 1.45 (Tychonoff)**.** *The arbitrary product of compact topological spaces is compact.*

Exercise 1.46. Look up a proof of Tychonoff's theorem.[a]

[a]On your quest you might encounter *ultrafilters*, which are one way to furnish a replacement for sequences in arbitrary topological spaces. They permit an elegant proof of Tychonoff's theorem, at the cost of some technical baggage (if you like, within which the core of the proof is really hidden).

Exercise 1.47. Famously, Tychonoff's theorem for arbitrary products requires the axiom of choice, and Tychonoff's theorem for merely countable products requires the ultrafilter lemma. Where did we use the axiom of choice in the proof of Theorem 1.43 in Subsection 1.5?

Next we establish an equivalent characterization of compactness which is sometimes convenient.

Definition 1.48. We say that a family $\mathcal{F}$ of subsets of a set S has the *finite intersection property (FIP)* if every finite subset of $\mathcal{F}$ has nonempty intersection.

Theorem 1.49. *A space X is compact if and only if every family $\mathcal{F}$ of closed subsets of X with the FIP has $\bigcap \mathcal{F} \neq \emptyset$.*

Proof. Observe that if $\mathcal{U}$ is an open cover of X then $\mathcal{F}_{\mathcal{U}} = \{U^c : U \in \mathcal{U}\}$ is a family of closed subsets of X satisfying $\bigcap \mathcal{F}_{\mathcal{U}} = \emptyset$. Indeed, it is easy to see that $\mathcal{U} \mapsto \mathcal{F}_{\mathcal{U}}$ is a one-to-one correspondence between open covers and families of closed sets with empty intersection.

We show the contrapositives: first suppose that X is compact and let $\mathcal{F}$ be a family of closed subsets of X with $\bigcap \mathcal{F} = \emptyset$. Then $\mathcal{F} = \mathcal{F}_{\mathcal{U}}$ for an open cover $\mathcal{U}$ of X. There is a finite subcover $\mathcal{V}$ of $\mathcal{U}$, for which we necessarily have $\bigcap \mathcal{F}_{\mathcal{V}} = \emptyset$. Since $\mathcal{F}_{\mathcal{V}} \subseteq \mathcal{F}_{\mathcal{U}}$ it follows that $\mathcal{F} = \mathcal{F}_{\mathcal{U}}$ does not have the FIP, as desired. The other direction is analogous and is left as an exercise. □

Exercise 1.50. Finish the proof of Theorem 1.49.

The purpose of the remainder of this section is to introduce a certain finiteness condition on our spaces called "paracompactness", in order to state an important consequence called "existence of partitions of unity". Partitions of unity in turn allow us to define functions on an entire topological space X by specifying arbitrary functions on each element of an open cover of X. This is very powerful, and in class will e.g. allow us to define structures on our vector bundles (whatever vector bundles are) by chopping them up into manageable pieces specified by an open cover.

Definition 1.51. A *cover* $\mathcal{U}$ of X is *locally finite* if for each $x \in X$ there is an open set $V \subseteq X$ containing x such that V has nonempty intersection with only finitely many $U \in \mathcal{U}$.

Proposition 1.52. *Any open cover $\mathcal{U}$ of a compact space X has a locally finite open subcover.*

Proof. Since X is compact $\mathcal{U}$ admits a finite open subcover, and finite covers are locally finite. □

Example 1.53. Of course, in general there are plenty of locally finite open covers which do not admit finite subcovers. Consider for instance the open cover $\mathcal{U} = \{(n-1, n+1) : n \in \mathbb{Z}\}$ of $\mathbb{R}$.

Definition 1.54. A topological space X is *paracompact* if each open cover of X admits a locally finite refinement.

Example 1.55. Proposition 1.52 shows that compact spaces are paracompact.

Example 1.56. The space $\mathbb{R}$, and in general $\mathbb{R}^n$, is paracompact. In general, all metric spaces are paracompact (see Theorem 1.59 below). If you know what manifolds are: this implies that all (second-countable) manifolds are paracompact.[3]

Though we have avoided it so far, we now introduce the notion of a Hausdorff topological space. This is a certain very weak condition which all of the spaces we care about in algebraic topology satisfy, and which means that distinct points of the space may be separated by disjoint open sets.

Definition 1.57. A topological space X is *Hausdorff* if whenever $x, y \in X$ are distinct there exist disjoint open sets U_x and U_y such that $x \in U_x$ and $y \in U_y$.

Exercise 1.58. Show that a topological space X is Hausdorff if and only if the diagonal $\Delta = (x, x) : x \in X \subset X \times X$ is closed (in the product topology).

The Smirnov metrization theorem (which we will not use) uses paracompactness to characterize all of the metric spaces. Note that a space is *metrizable* if it is

[3]Something called the *long line* (a topological space which looks like $\mathbb{R}$ but just much, much longer) is an example of a non-second-countable manifold which is not paracompact.

homeomorphic to a metric space.

Theorem 1.59 (Smirnov)**.** *A space X is metrizable if and only if X is Hausdorff, paracompact, and locally metrizable*[a]*.*

Proof. This is Theorem 42.1 of Munkres [4]. □

[a] A space X is locally metrizable if every $x \in X$ is contained in an open set U which is metrizable (as a subspace of X).

As alluded to earlier, we care about paracompactness because we care about existence of so-called "partitions of unity". Roughly speaking, a partition of unity on a space X subordinate to a fixed open cover $\mathcal{U}$ is a way of writing the constant function 1 on X as a sum of nonnegative functions which are each guaranteed to be zero outside of a particular open set $U \in \mathcal{U}$. The precise definition follows.

Definition 1.60. If $\mathcal{U}$ is an open cover of X, then a *partition of unity (subordinate to $\mathcal{U}$)* is a family of functions $\{f_U : X \to [0,1]\}_{U \in \mathcal{U}}$ such that

1. (*local finiteness*) for each $x \in X$ there is an open set $V \subseteq X$ with $x \in V$ such that there are only finitely many $U \in \mathcal{U}$ for which $f_U|_V \neq 0$,
2. (*partition of unity*) in particular for each $x \in X$ the sum $\sum_{U \in \mathcal{U}} f_U(x)$ is finite, we require in addition that $\sum_{U \in \mathcal{U}} f_U(x) = 1$, and
3. (*subordinate to $\mathcal{U}$*) for each $U \in \mathcal{U}$ we have $\overline{f_U^{-1}((0,1])} \subseteq U$.

The second condition is the sense in which the family $\{f_U\}$ is a partition of unity[4], and the third condition says that the partition is subordinate to the cover $\mathcal{U}$. The first condition is a technical one which (though nonetheless usually part of the definition) ensures that partitions of unity will play nicely with the constructions on vector bundles which we will see in class. Actually, it turns out that every family of functions $\{f_U\}_{U \in \mathcal{U}}$ which is merely *point finite*[5] and satisfies (2) and (3) admits a refinement in a certain sense which is locally finite in the sense of condition (1) in the definition above, so the distinction is not really material—but the standard formulation is simplest for our purposes.[6].

Theorem 1.61. *Let X be a Hausdorff space. Then X is paracompact if and only if every open cover $\mathcal{U}$ of X admits a partition of unity subordinate to $\mathcal{U}$.*

The remainder of this section is devoted to sketching the proof of Theorem 1.61: this again requires far more point-set topology than we will need in class, and we'll hide some of the details in a pair of citations inside the proof of Proposition 1.63.

[4] The term "unity" is an old-fashioned name for the number 1.

[5] We say that the family $\{f_U\}_{U \in \mathcal{U}}$ is *point finite* if for each $x \in X$ we have $f_U(x) \neq 0$ for only finitely many $U \in \mathcal{U}$.

[6] This is a result of Mather (1965): a proof appears as e.g. Proposition A.2.8 of [1].

We first need one more technical definition.

Definition 1.62. Let $\mathcal{U}$ be an open cover of a topological space X. A *shrinking* of $\mathcal{U}$ is another open cover $\mathcal{V} = \{V_U\}_{U \in \mathcal{U}}$ indexed by the elements of $\mathcal{U}$ such that $\overline{V_U} \subseteq U$ for all $U \in \mathcal{U}$.

A topological space X is *shrinking* if every open cover of X admits a *shrinking*.

Proposition 1.63. *Let X be a paracompact Hausdorff space. If A and B are any two disjoint closed subsets of X then there exists a continuous function $f : X \to [0,1]$ such that $f|_A = 1$ and $f|_B = 0$.*

Proof. Theorem 41.1 of Munkres [4] asserts that paracompact Hausdorff spaces are so-called *normal* spaces. Then the Urysohn lemma (a basic, deep theorem in point-set topology, see Theorem 33.1 of Munkres [4]) asserts that such a functions $f : X \to [0,1]$ exist for all normal spaces. □

Proposition 1.64. *Every paracompact Hausdorff space X is shrinking.*

Proof. Let $\mathcal{U}$ be an arbitrary open cover of X. Define

$$\mathcal{W} = \{W \text{ open} : \overline{W} \subseteq U \text{ for some } U \in \mathcal{U}\}$$

and observe that by Proposition 1.63 $\mathcal{W}$ is actually a refinement of $\mathcal{U}$ (exercise). By paracompactness of X then find a locally finite subcover $\mathcal{W}'$ of $\mathcal{W}$. For each $W \in \mathcal{W}'$ fix a specific $U_W \in \mathcal{U}$ such that $\overline{W} \subseteq U_W$ (there must be at least one such U_W for each W). Also, for each $U \in \mathcal{U}$ now define

$$V_U = \bigcup \{W \in \mathcal{W}' : U = U_W\}.$$

By construction, the family $\mathcal{V} = \{V_U\}_{U \in \mathcal{U}}$ still covers X and is still a refinement of $\mathcal{U}$ (in the latter case because $V_U \subseteq U$). Since the cover $\mathcal{W}'$ is locally finite the closure $\overline{V_U}$ may be calculated as the the union of the closures of sets $W \in \mathcal{W}'$ with $U = U_W$ (exercise), and thus $\overline{V_U} \subseteq U$ for all $U \in \mathcal{U}$. Therefore $\mathcal{V}$ is the desired refinement of $\mathcal{U}$. □

Proof of Theorem 1.61. One direction is straightforward: first suppose that X admits partitions of unity subordinate to all open covers. Then if $\mathcal{U}$ is a fixed open cover, there is a partition of unity $\{f_U\}_{U \in \mathcal{U}}$ subordinate to U. Each preimage $V_U := f_U^{-1}((0,1])$ is an open subset of U, and each $x \in X$ is contained in V_U for some $U \in \mathcal{U}$ because $\sum_{U \in \mathcal{U}} f_U(x) = 1 \neq 0$. Therefore $\mathcal{V} = \{V_U : U \in \mathcal{U}\}$ is a refinement of $\mathcal{U}$, which is locally finite by condition (1) in the definition of a partition of unity (Definition 1.60).

Now instead suppose that X is paracompact and let $\mathcal{U}$ be an arbitrary open cover of X. First observe that it suffices to assume that $\mathcal{U}$ is itself locally finite (exercise). Assuming this, then by Proposition 1.64 X is shrinking, so there is a shrinking $\mathcal{V} = \{V_U\}_{U \in \mathcal{U}}$ of $\mathcal{U}$ and a further shrinking $\mathcal{W} = \{W_U\}_{U \in \mathcal{U}}$ of $\mathcal{V}$ (with both $\mathcal{V}$ and $\mathcal{W}$ necessarily locally finite). By Proposition 1.63 for each $U \in \mathcal{U}$ there exists a continuous function $f_U : X \to [0,1]$ such that $f_U|_{\overline{W_U}} = 1$ and $f_U|V_U^c = 0$. This implies property (3) of Definition 1.60 that $\overline{f_U^{-1}((0,1])} \subseteq U$ for each $U \in \mathcal{U}$ (exercise). Moreover, the fact that the cover $\mathcal{V}$ is locally finite implies property (1) of Definition 1.60 (exercise). Finally, for each $x \in X$ we may define a function $f : X \to [0,1]$ by

$$f(x) = \sum_{U \in \mathcal{U}} f_U(x)$$

since property (1) implies that this is actually a finite sum for each fixed $x \in X$. The resulting function f is continuous (exercise), and is strictly positive since each $x \in X$ belongs to W_U for at least one $U \in \mathcal{U}$. The functions $\{\frac{f_U}{f}\}_{U \in \mathcal{U}}$ then satisfy property (2) of Definition 1.60 while retaining properties (1) and (3), and therefore form the desired partition of unity. □

Exercise 1.65. Resolve all of the claims deferred as exercises in the proofs of Proposition 1.64 and Theorem 1.61. If you get stuck, look up Theorem 41.7 and the preceding Lemma 41.6 of Munkres [4].

1.4 Homotopy theory

In this tutorial we won't need much homotopy theory (though this is a source of many good ideas for final papers!). We will however want to borrow a few basic notions, because they appear in something called the "classification theorem" for real and complex vector bundles. The most basic definition is the following.

Definition 1.66. Let X and Y be topological spaces. A *homotopy* is a continuous map $I \times X \to Y$.[a] Two maps $f_0, f_1 : X \to Y$ are *homotopic* if there exists a homotopy $F : I \times X \to Y$ such that $f_0(x) = F(0,x)$ and $f_1(x) = F(1,x)$.

[a]Recall that $I = [0,1]$ is the unit interval.

Remark 1.67. Fix a homotopy F from f_0 to f_1. We can view t as a "time parameter", and for each fixed time $t \in I$ observe that the homotopy F gives us a map $f_t(x) = F(t,x) : X \to Y$. Thus a homotopy between f_0 and f_1 can be thought of as a movie which continuously deforms one map f_0 into another map f_1.

Exercise 1.68. If $x_0, x_1 \in X$, a *path* from x_0 to x_1 is a continuous map $f : I \to X$ such that $f(0) = x_0$ and $f(1) = x_1$. For each $x \in X$ let $\underline{x}$ denote the inclusion $\bullet \hookrightarrow X$ of the one-point space $\bullet$ into X with image $\{x\}$. Show that a path from x_0 to x_1 is the same a a homotopy between $\underline{x_0}$ and $\underline{x_1}$. Also show that there is a path between any two points in (1) $\mathbb{R}^n$, and (2) the unit circle S^1.[a]

[a]A space which has paths between each pair of points is called *path-connected*.

Definition 1.69. A map $f : X \to Y$ is a *homotopy equivalence* if there exists another map $g : Y \to X$ such that $g \circ f$ and $f \circ g$ are each homotopic to the identity map (on X and Y respectively). Such a map g is called a *homotopy inverse* of f.

Definition 1.70. A topological space X is *contractible* if X is homotopy equivalent to the the one-point space $\bullet$.

Exercise 1.71. Show that homotopy equivalence is an equivalence relation on the set[a] of all topological spaces.

[a] In truth the collection of all topological spaces is too large to be set, but we won't dwell on the way to fix this and no harm is done in adopting the naive interpretation.

Exercise 1.72. Show that a topological space X is contractible if and only if the identity map $X \to X$ is homotopic to a constant map.

Exercise 1.73. Show that $\mathbb{R}$ and I are both contractible. On the other hand, the unit circle $S^1 \subset \mathbb{R}^2$ is not contractible, but it is a little bit above our pay-grade to see this here: look up the proof, which is e.g. Theorem 1.7 of Hatcher's *Algebraic Topology* [3].

Finally, in preparation for the classification theorem for vector bundles which we'll see in class, we need one more definition.

Definition 1.74. Let X and Y be topological spaces. Then the set $[X \to Y]$ of *homotopy classes of maps* $X \to Y$ is the set of all continuous functions $X \to Y$ modulo the equivalence relation that $f \sim g$ whenever f and g are homotopic.

Exercise 1.75. Show that the relation used to define $[X \to Y]$ is actually an equivalence relation.

Exercise 1.76. Let X, Y, and Z all be topological spaces.

1. Show that the sets $[X \sqcup Y \to Z]$ and $[X \to Z] \sqcup [Y \to Z]$ are in canonical bijection.

2. Show that the sets $[X \to Y \times Z]$ and $[X \to Y] \times [X \to Z]$ are in canonical bijection.

3. Show that if Y is contractible then the set $[X \to Y]$ has one element.

1.5 Compact metric spaces

This subsection is devoted to proving Theorem 1.43. It's all entirely optional, and we won't even be caring specifically about metric spaces in class—it just might be interesting if you like metric spaces. If you don't like the sound of this, skip to Section 2 on linear algebra.

We will also invoke some more advanced notions in this subsection (like first/second-countability, complete/totally bounded metric spaces, and Lindelöf spaces), which we haven't defined earlier and which you won't remotely need to know the definitions of for our class. So, don't fret if you haven't heard of any or all of these (e.g. I doubt my fellow graduate students all remember the definition of a Lindelöf space).

Before proving Theorem 1.43 it will be useful to introduce the following intermediary weakenings of compactness.

Definition 1.77. A space X is *Lindelöf* if every open cover has a countable subcover, and is *countably compact* if every countable open cover has a finite subcover.

The proof of Theorem 1.43 mainly amounts to establishing the following chain of implications for metric spaces (X, d):

$$\begin{aligned} X \text{ is compact} &\implies X \text{ is sequentially compact} \\ &\implies X \text{ is totally bounded} \\ &\implies X \text{ is separable} \\ &\implies X \text{ is second-countable} \\ &\implies X \text{ is Lindelöf.} \end{aligned}$$

We will also show that all sequentially compact topological spaces (not necessarily metric spaces) are countably compact. Since countably compact Lindelöf spaces are compact (just by definition), this then completes the proof. Let's now handle each of these implications in turn.

Proposition 1.78. *Every sequentially compact topological space is countably compact.*

Proof. We show the contrapositive: let $\mathcal{U} = \{U_n\}_{n\in\mathbb{N}}$ be a countable open cover of a topological space X with no finite subcover. In particular, the set $C_n = X \setminus \bigcup_{k=1}^{n} U_k$ is nonempty for all $n \in \mathbb{N}$. Now define a sequence (x_n) in X by choosing some $x_n \in C_n$ for each $n \in \mathbb{N}$. We claim that (x_n) has no convergent subsequence: in order to verify this fix $x \in X$. Because $X = \bigcup_{n=1}^{\infty} U_n$ there is some $N \in \mathbb{N}$ for which $x \in U_N$, and in particular $x_n \notin U_N$ for any $n > N$. Therefore there is no subsequence of (x_n) which converges to x, as desired. □

Proposition 1.79. *Compact metric spaces are sequentially compact.*

Proof. We again show the contrapositive: suppose that the metric space (X, d) is not sequentially compact, and let (x_n) be a sequence in X with no convergent subsequence. This means that for each fixed $x \in X$ there exists $\varepsilon_x > 0$ and $N_x \in \mathbb{N}$ so that $d(x_n, x) \geq \varepsilon_x$ for all $n \geq N_x$. The set $\mathcal{U}$ of all open balls $U_x = B(x, \varepsilon_x)$ of radius ε_x about each fixed $x \in X$ is then an open cover of X. However, this open cover $\mathcal{U}$ has no finite subcover $\mathcal{V}$: let $\mathcal{V} = \{U_{x_1}, \ldots, U_{x_n}\} \subseteq \mathcal{U}$ be a finite subset. Setting $N = \max\{N_{x_1}, \ldots, N_{x_n}\}$ we find that x_N is not in $\bigcup \mathcal{V}$, and therefore $\mathcal{V}$ is not a cover, as desired. □

Exercise 1.80. In fact Proposition 1.79 is a special case of the general fact that countably compact first-countable spaces are sequentially compact. That is, the converse of Proposition 1.78 holds for first-countable spaces. The proof is more difficult, and one possible strategy proceeds as follows:

1. First show that if x is an accumulation point of a sequence (x_n) in an arbitrary first-countable topological space X, then there is a subsequence which converges to x.

2. Complete the proof by showing that all infinite sequences in countably compact spaces have an accumulation point. In particular, suppose that (x_n) is a sequence in X with no accumulation point, and then:

   (a) Show that for each $x \in X$ there is an open set U_x containing x such that the set $S_x = \{n \in \mathbb{N} : x_n \in U_x\}$ is finite.

   (b) For each finite subset $S \subset \mathbb{N}$ define $U_S = \bigcup\{U_x : x \in X, S_x = S\}$, and show that $\mathcal{U} = \{U_S : S \subset \mathbb{N} \text{ is finite}\}$ is a countable open cover of X.

   (c) Show that if $\mathcal{V} = \{U_{S_1}, \ldots, U_{S_n}\} \subset \mathcal{U}$ is any finite subset then $\mathcal{V}$ is not a cover of X.

   (d) Conclude that $\mathcal{U}$ has no a finite subcover, as desired.

Theorem 1.81. *A metric space is sequentially compact if and only if it is complete and totally bounded.*

Proof. Let (X, d) be a metric space. First note that if X is not complete then there is a Cauchy sequence (x_n) in X which does not converge. Then no subsequence of (x_n) converges either, so X is not sequentially compact. If instead X is not totally bounded then there is some $\varepsilon > 0$ for which X does not admit a finite cover by ε-balls. We can then inductively build a sequence of points (x_n) of X starting by letting $x_1 \in X$ be arbitrary. Then at the $(k+1)$st stage choose $x_{k+1} \in X \setminus \bigcup_{j=1}^{k} B(x_j, \varepsilon)$, and note that we always have $\bigcup_{j=1}^{k} B(x_j, \varepsilon) \neq X$ because by hypothesis X does not admit a finite cover by ε-balls. By construction the infinite sequence (x_n) satisfies $d(x_m, x_n) \geq \varepsilon$ for all $m \neq n$, and therefore has no convergent subsequence. Hence X is not sequentially compact.

Finally, now assume that X is complete and totally bounded, and fix any sequence (x_n) in X. If the set $\{x_n\}$ is finite then X has a constant (hence convergent) subsequence, so also assume that $\{x_n\}$ is infinite. Because X is complete it suffices to produce a Cauchy subsequence of (x_n), which we will now construct inductively. First let $(x_{j_{0,n}})$ be the subsequence of (x_n) equal to (x_n) itself. At the $(k+1)$th stage cover X by finitely many balls $B_1, \ldots, B_N$ of radius $\frac{1}{k+1}$. Since we may assume that the set $\{x_{j_{k,n}} : j \in \mathbb{N}\}$ is infinite there must be some ball B_i for which the intersection $B_i \cap \{x_{j_{k,n}} : j \in \mathbb{N}\}$ is infinite. Now let $(x_{j_{k+1,n}})$ be a subsequence of $(x_{j_{k,n}})$ which is wholly contained in B_i.

This process produces an infinite k-indexed family of further subsequences $(x_{j_{k,n}})$ of the original sequence (x_n), and which all satisfy the property that for each fixed $K \in \mathbb{N}$ every sequence $(x_{j_{k,n}})$ with $k \geq K$ is wholly contained in the same fixed ball of radius $\frac{1}{K}$. In particular, we are guaranteed that the sequence (y_n) defined by $y_n = x_{j_{n,n}}$ is a subsequence of (x_n), and by construction each set $\{y_n : n \geq N\}$ is contained in a ball of radius $\frac{1}{N}$. This forces the sequence (y_n) to be a Cauchy subsequence of the sequence (x_n), as desired. □

Proposition 1.82. *Totally bounded metric spaces are separable.*

Proof. If (X, d) is totally bounded then for each fixed $n \in \mathbb{N}$ there is a finite subset $S_n \subseteq X$ such that $X = \bigcup_{x \in S_n} B(x, \frac{1}{n})$. The union $S = \bigcup_{n \in \mathbb{N}} S_n$ is a countable set. Moreover S is dense in X because $d(x, S) \leq d(x, S_n) < \frac{1}{n}$ for all $n \in \mathbb{N}$, and therefore $d(x, S) = 0$. □

Proposition 1.83. *Second-countable topological spaces are Lindelöf.*

Proof. Let $\mathcal{U}$ be an open cover of a topological space X. By hypothesis X is second-countable with countable base $\mathcal{B}$. Define a subset of $\mathcal{B}$ by

$$\mathcal{B}' = \{B \in \mathcal{B} : \exists U \in \mathcal{U}, B \subseteq U\}.$$

Because $\mathcal{B}$ is a base we must have that $U \subseteq \bigcup \mathcal{B}'$ for all $U \in \mathcal{U}$, so $\mathcal{B}'$ is a countable open cover of X. For each $B \in \mathcal{B}'$ choose a particular $U_B \in \mathcal{U}$ such that $B \subseteq U_B$. Then $\mathcal{U}' = \{U_B : B \in \mathcal{B}'\}$ is a countable subcover of $\mathcal{U}$, as desired. □

Proof of Theorem 1.43. If X is compact then Proposition 1.79 yields that X is sequentially compact. On the other hand if X is sequentially compact then X is totally bounded by Theorem 1.81, hence separable by Proposition 1.82. Since metric spaces are always first-countable it follows that X is second-countable, hence X is Lindelöf by Proposition 1.83. By appeal to Proposition 1.78 we conclude that X is countably compact Lindelöf, and therefore compact, as desired. □

2 Linear algebra

On our first day we'll define vector bundles, which are in a sense generalized vector spaces. We'll find that all of the operations we can perform on vector spaces can be performed on vector bundles, so we're going to take this opportunity to review and catalog a bunch of constructions for ordinary vector spaces.

Note that for vector spaces some of these are so simple they are a bit silly, but I promise that they all give rise to interesting operations at the level of vector bundles.

To start with all of our vector spaces will be over an arbitrary field $\Bbbk$ of characteristic zero (but if this worries you then you can assume $\Bbbk = \mathbb{R}$ or $\mathbb{C}$).

2.1 Homs, duals, and direct sums

The first natural constructions come from considering maps between vector spaces.

Definition 2.1. If U and V are vector spaces then the *hom-space* $\mathcal{L}(U \to V)$ of linear maps from U to V is again a vector space.

Exercise 2.2. Show that if U and V are finite-dimensional then $\mathcal{L}(U \to V)$ is too. What is the dimension of $\mathcal{L}(U \to V)$?

Note that we can always view $\Bbbk$ as a 1-dimensional vector space over itself.

Definition 2.3. The *dual* vector space V^* is the hom-space $\mathcal{L}(V \to \Bbbk)$.

Exercise 2.4. Let V be a finite dimensional vector space. Show that $\dim V^* = \dim V$, and conclude that V and V^* are (non-canonically) isomorphic. In general there is no preferred choice of isomorphism $V \to V^*$, and such isomorphisms are closely related[a] to the theory of inner products on V.

[a]The interested reader can consult Theorem 2.49 and the proof immediately preceding it. Isomorphisms $V \to V^*$ are often called *nondegenerate bilinear forms*.

Exercise 2.5. Despite the caveat in Exercise 2.4 that isomorphisms $V \cong V^*$ are usually non-canonical, show that in the special case that L is a one-dimensional vector space then the hom-space $\mathcal{L}(L \to L)$ is canonically isomorphic to $\Bbbk$. Hence conclude that $\Bbbk^* = \mathcal{L}(\Bbbk \to \Bbbk)$ is canonically isomorphic to $\Bbbk$.

Hint: there are several ways to recognize this isomorphism. One way is via the trace, but a simpler method is just to observe that every element of $\mathcal{L}(L \to L)$ is uniquely a scalar multiple of the identity map $L \to L$.

In all of the constructions to come we will find (in Subsection 2.5) that linear maps (e.g. $U_1 \to V_1$ and $U_2 \to V_2$) will naturally induce linear maps between the objects we construct (e.g. $U_1 \oplus U_2 \to V_1 \oplus V_2$).[7] The next exercise explains how linear maps induce maps between dual spaces.

[7]This is because all of these constructions are all examples in one way or another of *functors*, a concept in *category theory*.

Exercise 2.6. Show that if $f : U \to V$ is a linear map then the function $f^* : V^* \to U^*$ defined by $\psi \mapsto \psi \circ f$ is a linear map. The map f^* is called the *dual* or sometimes *transpose* of f.

In particular conclude that if $U \subseteq V$ is a subspace and $\iota_U : U \hookrightarrow V$ is the inclusion then there is a canonical surjective *restriction* map $\iota_U^* : V^* \twoheadrightarrow U^*$.

Exercise 2.7. An element α of the *double dual* V^{**} is by definition a linear map $V^* \to \Bbbk$, which in particular assigns an element of $\Bbbk$ to each dual in V^*. For each $v \in V$ let α_v be the element of V^{**} which is *evaluation-at-v*, i.e. which maps each $\psi \in V^*$ to $\psi(v) \in \Bbbk$.

1. Show that the assignment $v \mapsto \alpha_v$ defines a linear injection $J_V : V \to V^{**}$.
2. Show that V is finite dimensional if and only if $J_V : V \to V^{**}$ is an isomorphism.
3. Conclude that, when V is finite dimensional, V and V^{**} are canonically isomorphic even though V and V^* need not be in general.
4. Show that the isomorphism J_V is *natural*[a], in the sense that for all linear maps $f : V \to W$ (between vector spaces of any dimension) we have $J_W \circ f = f^{**} \circ J_V$.[b]

[a] If you know a bit of category theory, in this exercise we're showing that the linear maps J_V define a natural transformation from the identity functor on the category of vector spaces to the the double dual functor. When we restrict to the full subcategory of finite-dimensional vector spaces, we've also shown above that this natural transformation becomes an isomorphism.

[b] The double dual map f^{**} is defined by taking the dual of the map $f : V \to W$ once to get a map $f^* : W^* \to V^*$, and again to finally obtain $f^{**} : V^{**} \to W^{**}$ (with the dual of a linear map defined as in Exercise 2.6).

Definition 2.8. If U and V are vector spaces then the *direct sum* $U \oplus V$ is another vector space which has underlying set $U \times V$. We define addition and scalar multiplication by

$$(u_1, v_1) + (u_2, v_2) = (u_1 + u_2, v_1 + v_2) \quad \text{and} \quad \lambda \cdot (u, v) = (\lambda u, \lambda v).$$

We write $u \oplus v$ for (u, v).

Exercise 2.9. What is the dimension of $U \oplus V$ in terms of $\dim U$ and $\dim V$?

2.2 Tensor products

Just as the direct sum of vector spaces $U \oplus V$ gives a precise meaning to formal sums $u \oplus v$ of vectors $u \in U$ and $v \in V$, the tensor product $U \otimes V$ is meant to give a precise meaning to formal products $u \otimes v$. The distinction is that, for formal sums we have

$$(u_1 \oplus v_1) + (u_2 \oplus v_2) = (u_1 + u_2) \oplus (v_1 + v_2) \quad \text{and} \quad \lambda \cdot (u \oplus v) = \lambda u \oplus \lambda v,$$

meanwhile formal products should satisfy

$$(u_1 + u_2) \otimes (v_1 + v_2) = (u_1 \otimes v_1) + (u_1 \otimes v_2) + (u_2 \otimes v_1) + (u_2 \otimes v_2) \tag{2.1}$$

and

$$\lambda \cdot (u \otimes v) = \lambda u \otimes v = u \otimes \lambda v. \tag{2.2}$$

One way to rigorously construct $U \otimes V$ proceeds by essentially just declaring (2.1) and (2.2) to be true: first, form the free $\Bbbk$-vector space $A = \Bbbk[\{u \otimes v : u \in U, v \in V\}]$ on all formal products $u \otimes v$. There are two things to note here:

1. At this point the formal products $u \otimes v$ are strictly just symbols which we can multiply by scalars, add, and subtract in A. It isn't yet the case that e.g. $2(u \otimes v) = 2u \otimes v = u \otimes 2v$, and likewise for the more complicated distributive law (2.1).

2. The free $\Bbbk$-vector space A is "very large": it is uncountably-infinite dimensional whenever $\Bbbk$ is uncountable and $\dim U, \dim V > 0$.

We can now define $U \otimes V$ by declaring (2.1) and (2.2) to be true: formally, we take the quotient of A by the subspace R generated by all elements of the form (for any $u_1, u_2 \in U$ and $v_1, v_2 \in V$)

$$(u_1 + u_2) \otimes (v_1 + v_2) - (u_1 \otimes v_1) - (u_1 \otimes v_2) - (u_2 \otimes v_1) - (u_2 \otimes v_2),$$

and (for any $\lambda \in \Bbbk$, $u \in U$, and $v \in V$)

$$\lambda(u \otimes v) - \lambda u \otimes v \quad \text{and} \quad \lambda(u \otimes v) - u \otimes \lambda v.$$

The subspace R of A generated by these relations is also "very large" in the informal sense used above, and so the quotient A/R collapses many elements (as Exercise 2.12 below shows, in fact whenever U and V are both finite dimensional then $U \otimes V$ is, too).

Definition 2.10. The *tensor product* $U \otimes V$ of vector spaces is the quotient A/R.

Sometimes people refer to elements of $U \otimes V$ as *tensors*.

Remark 2.11. Danger! Note that while in the direct sum $U \oplus V$ arbitrary finite sums

$$(u_1 \oplus v_1) \oplus \cdots \oplus (u_n \oplus v_n)$$

can always be written uniquely as a single formal sum $(u_1 + \cdots + u_n) \oplus (v_1 + \cdots + v_n)$, this is not the case for the tensor product $U \otimes V$. That is, the sum

$$(u_1 \otimes v_1) + (u_2 \otimes v_2)$$

need not be equal a single $u \otimes v$ for any $u \in U$ and $v \in V$. Some people like to distinguish the elements of $U \otimes V$ which can be written as $u \otimes v$ by calling them the *pure tensors*.

That said, to define a linear map T from $U \otimes V$ to any other vector space W it suffices to specify what happens to only the pure tensors $u \otimes v$, since then the value of T on any sum of pure tensors is clear (we "extend linearly"). In other words, if we have a definition of $T(u \otimes v)$ for all $u \in U$ and $v \in V$, then implicitly we also know

$$T\left((u_1 \otimes v_1) + \cdots + (u_n \otimes v_n)\right) = T(u_1 \otimes v_1) + \cdots + T(u_n \otimes v_n).$$

We will discuss defining maps out of tensor products more in the coming exercises, and further in Subsection 2.5 below.

Exercise 2.12. Let U and V be finite dimensional $\Bbbk$-vector spaces and suppose that $\{u_1, \ldots, u_m\}$ and $\{v_1, \ldots, v_n\}$ are bases of U and V respectively.

- Use the relations (2.1) and (2.2) to show that the set $\mathcal{B} = \{u_i \otimes v_j : 1 \leq i \leq m, 1 \leq j \leq n\}$ spans $U \otimes V$. (In this problem we write $u \otimes v$ for both the associated purely formal product in A, and the image of this formal product in the quotient $A/R = U \otimes V$.)
- For each $1 \leq i \leq m$ and $1 \leq j \leq n$ show that the linear map $f_{i,j} : A \to \Bbbk$ defined by

  $$u_k \otimes v_l \mapsto \begin{cases} 1 & (k,l) = (i,j) \\ 0 & \text{otherwise} \end{cases}$$

  respects the relations (2.1) and (2.2) (i.e. is identically zero on the subspace $R \subset A$) and therefore descends to a map $\widetilde{f_{i,j}} : U \otimes V = A/R \to \Bbbk$ from the quotient.
- Using the fact that image of $u_i \otimes v_j$ under $f_{k,l}$ is nonzero if and only if $(k,l) = (i,j)$, conclude that $\mathcal{B}$ is a linearly independent subset of $U \otimes V$, and is therefore a basis of $U \otimes V$.
- Moreover, conclude that even though the free $\Bbbk$-vector space A we used to define $U \otimes V$ was very large, in fact its quotient $U \otimes V$ is finite dimensional and $\dim U \otimes V = \dim U \cdot \dim V$.

Exercise 2.13. Show that $\Bbbk$ is a unit for the tensor product, in the sense that given any (possibly infinite dimensional) vector space V, there are canonical isomorphisms

$$V \otimes \Bbbk \cong V \cong \Bbbk \otimes V.$$

Note that here we are again viewing $\Bbbk$ as a 1-dimensional $\Bbbk$-vector space.

Exercise 2.14. Show that the operation of evaluating $\psi \in V^*$ on a vector $v \in V$ gives rise to a canonical *evaluation* map $V^* \otimes V \to \Bbbk$.

Exercise 2.15. Show that when U is finite dimensional there is a canonical isomorphism

$$\mathcal{L}(U \to V) \cong V \otimes U^*.$$

By considering the case when $U = V$ conclude that when V is finite dimensional there is a canonical (basis-independent) nonzero element of $V \otimes V^*$, given by the image of the identity map $V \to V$ under this isomorphism. The linear map $\Bbbk \to V \otimes V^*$ which sends $1 \in \Bbbk$ to this element is called the *coevaluation* map.

2.3 Tensor, symmetric, and exterior powers

As a special case of the tensor product of vector spaces, we can always form the tensor product $V \otimes V$ of a vector space with itself. The next exercise shows that even though in this case the pure tensors $v_1 \otimes v_2$ and $v_2 \otimes v_1$ both lie in the same vector space $V \otimes V$, in general they need not be related.

Exercise 2.16. Show that the elements $v_1 \otimes v_2$ and $v_2 \otimes v_1$ of $V \otimes V$ are linearly dependent if and only if v_1 and v_2 are linearly dependent. Hint: use the explicit basis of Exercise 2.12.

Of course, there is no reason to limit ourselves to only two copies of V.

Definition 2.17. For any $k \geq 1$ we can form the k-fold tensor product of V with itself, obtaining the *kth tensor power*

$$V^{\otimes k} := \underbrace{V \otimes \cdots \otimes V}_{k\text{-times}}.$$

Since $\Bbbk$ is the unit for the tensor product we define $V^{\otimes 0} = \Bbbk$.

Exercise 2.18. Let $\{v_1, \ldots, v_n\}$ be a basis of V. By Exercise 2.12 the set $\{v_i \otimes v_j : 1 \leq i, j \leq n\}$ is a basis of $V \otimes V$. Generalize this and show that

$$\{v_{j_1} \otimes \cdots \otimes v_{j_k} : 1 \leq j_1, \ldots, j_k \leq n\}$$

of all possible k-fold products of these basis vectors is a basis of of $V^{\otimes k}$. Conclude that $\dim V^{\otimes k} = (\dim V)^k$. Hint: no work is required, this is a formal consequence of the previous exercise.

It is sometimes desirable to impose additional relations on formal k-fold products $v_1 \otimes \cdots \otimes v_k$ in $V^{\otimes k}$. By far the most common two are:

- Declare that the formal multiplication "$\otimes$" is commutative, in that

  $$v_1 \otimes \cdots \otimes v_j \otimes v_{j+1} \otimes \cdots \otimes v_k = v_1 \otimes \cdots \otimes v_{j+1} \otimes v_j \otimes \cdots \otimes v_k$$

  for all j.

- Declare that the formal multiplication "⊗" is anticommutative, in that

$$v_1 \otimes \cdots \otimes v_j \otimes v_{j+1} \otimes \cdots \otimes v_k = -(v_1 \otimes \cdots \otimes v_{j+1} \otimes v_j \otimes \cdots \otimes v_k) \tag{2.3}$$

for all j.

One formally imposes either of these relations on $V^{\otimes k}$ by taking a quotient, in a similar fashion to how we constructed $V^{\otimes k}$ itself.

Definition 2.19. The *kth symmetric power* $\mathrm{S}^k V$ is the quotient of $V^{\otimes k}$ by the subspace spanned by all elements of the form

$$v_1 \otimes \cdots \otimes v_j \otimes v_{j+1} \otimes \cdots \otimes v_k - v_1 \otimes \cdots \otimes v_{j+1} \otimes v_j \otimes \cdots \otimes v_k.$$

The *kth exterior power* $\Lambda^k V$ is the quotient of $V^{\otimes k}$ by the subspace spanned by all elements of the form

$$v_1 \otimes \cdots \otimes v_j \otimes v_{j+1} \otimes \cdots \otimes v_k + v_1 \otimes \cdots \otimes v_{j+1} \otimes v_j \otimes \cdots \otimes v_k.$$

In order to prevent confusion the image of $v_1 \otimes \cdots \otimes v_k \in V^{\otimes k}$ in the quotient defining $\mathrm{S}^k V$ is denoted $v_1 \circledS \cdots \circledS v_k$. Likewise, the image of $v_1 \otimes \cdots \otimes v_k$ in the quotient in the quotient defining $\Lambda^k V$ is denoted $v_1 \wedge \cdots \wedge v_k$.

Exercise 2.20. Let $\{v_1, \ldots, v_n\}$ be a basis of V. Show that

$$\{v_{j_1} \circledS \cdots \circledS v_{j_k} : 1 \le j_1 \le \cdots \le j_k \le n\}$$

of all possible k-fold products of basis vectors with **non-strictly** increasing index is a basis of $\mathrm{S}^k V$. Similarly, show that

$$\{v_{j_1} \wedge \cdots \wedge v_{j_k} : 1 \le j_1 < \cdots < j_k \le n\}$$

of all possible k-fold products of basis vectors with **strictly** increasing index is a basis of $\Lambda^k V$. Conclude that $\dim \mathrm{S}^k V = \binom{n+k-1}{k}$ and $\dim \Lambda^k V = \binom{n}{k}$.

Exercise 2.21. Let V be an n-dimensional vector space. Show that $(\Lambda^k V)^*$ is canonically isomorphic to $(\Lambda^{n-k} V)^* \otimes \Lambda^n V$.

Exercise 2.22. Use the fact that $\Bbbk$ is a field of characteristic zero to show that the kth exterior power $\Lambda^k V$ may equivalently be defined as the quotient of $V^{\otimes k}$ by the subspace spanned by all elements of the form

$$v_1 \otimes \cdots \otimes v_j \otimes v \otimes v \otimes v_{j+3} \otimes \cdots \otimes v_k,$$

i.e. where two adjacent vectors in the formal product are equal.

Note that over a field of characteristic 2 these two definitions of $\Lambda^k V$ are not the same: indeed, when $\operatorname{char} \Bbbk = 2$ then (as we have defined it) the kth exterior power becomes the same as the kth symmetric power, while this other definition in general gives a distinct object.

We will mostly be interested in symmetric and exterior powers for their own sake, since they directly give rise to corresponding operations on vector bundles. However, it is worth mentioning that in particular exterior powers of vector bundles naturally appear in the integration theory of differential geometry because the relation (2.3) is fundamentally related to measurement of oriented areas and volumes.

Another application of exterior powers is as an easy and intrinsic way to define the determinant of a linear map $V \to V$ (along with assorted machinery, such as the adjugate) without having to choose a basis. Going into this now would take us too far afield, but I'd be happy to tell you more about this story if you ask me. Coordinate-free linear algebra is fun!

2.4 Real and complex vector spaces

Up until now we have been dealing with arbitrary $\Bbbk$-vector spaces; we'll now assume that our vector spaces are real or complex. The purpose of this subsection is to enumerate some miscellaneous special features of the real or complex case. First, just as we can take the complex conjugate of a complex number, we can take the complex conjugate of a $\mathbb{C}$-vector space.

Definition 2.23. If V is a $\mathbb{C}$-vector space then the *complex conjugate* $\overline{V}$ is another $\mathbb{C}$-vector space which has the same underlying set and sum operation as V, but where we redefine the scalar multiplication to be $\lambda \cdot v := \overline{\lambda} v$.

We also have a general way to pass between the worlds of real and complex vector spaces. One direction comes from the fact that if V is a $\mathbb{C}$-vector space, then since V has scalar multiplication by complex numbers, V certainly has scalar multiplication by real numbers.

Definition 2.24. Let V be a $\mathbb{C}$-vector space. The *realification* $V_{\mathbb{R}}$ of V is the $\mathbb{R}$-vector space with the same underlying set as V, but where we forget the multiplication by complex scalars in V and remember only the multiplication by real scalars.

Exercise 2.25. Let V be a finite dimensional complex vector space. Then $\dim_{\mathbb{C}} V < \infty$, where we have written a $\mathbb{C}$ subscript to emphasize that we are viewing V as a $\mathbb{C}$-vector space. Similarly the real vector space $V_{\mathbb{R}}$ has dimension $\dim_{\mathbb{R}} V_{\mathbb{R}}$. Show that $\dim_{\mathbb{R}} V_{\mathbb{R}} = 2 \dim_{\mathbb{C}} V$.

Now suppose instead that we start with an $\mathbb{R}$-vector space V. As a special case of the realification construction, we can turn $\mathbb{C}$ into a real vector space $\mathbb{C}_{\mathbb{R}}$, which by Exercise 2.25 has (real) dimension 2. We can therefore take the tensor product of the real vector spaces $\mathbb{C}_{\mathbb{R}}$ and V.

Definition 2.26. Let V be an $\mathbb{R}$-vector space. The *complexification* $V_{\mathbb{C}}$ of V is the $\mathbb{C}$-vector space with the same underlying set as $\mathbb{C}_{\mathbb{R}} \otimes V$, but where we define the multiplication by an arbitrary scalar $\lambda \in \mathbb{C}$ as[a]

$$\lambda \cdot (\mu \otimes v) := \lambda\mu \otimes v.$$

[a] As always, this definition on pure tensors extends linearly to a definition for arbitrary sums of pure tensors.

Exercise 2.27. Show that $\mathbb{C}$-multiplication given in Definition 2.26 is compatible with addition, and hence that $V_{\mathbb{C}}$ is actually a $\mathbb{C}$-vector space.

Exercise 2.28. Show that the complexification obeys $\dim_{\mathbb{C}} V_{\mathbb{C}} = \dim_{\mathbb{R}} V$.

Exercise 2.29. If V is an $\mathbb{R}$-vector space then write down an explicit isomorphism $(V_{\mathbb{C}})_{\mathbb{R}} \cong V \oplus V$ of real vector spaces.

Recall that we also have the notion of an inner product on a real or complex vector space.

Definition 2.30. Let V be a $\Bbbk$-vector space with $\Bbbk = \mathbb{R}$ or $\mathbb{C}$. An *inner product* $(\cdot,\cdot)$ on V is a bilinear map[a] $V \times \overline{V} \to \Bbbk$ which is

- *conjugate symmetric*, in that[b] $(v_1, v_2) = \overline{(v_2, v_1)}$ for all $v_1, v_2 \in V$, and
- *positive semidefinite*, in that[c] $(v, v) \geq 0$ for all $v \in V$, and
- *definite*, in that $(v, v) = 0$ implies $v = 0$ for all $v \in V$.

[a] If V is real then we adopt the convention that $\overline{V} = V$. Recall that whenever U, V, and W are $\mathbb{R}$-vector spaces a *bilinear* map $f : U \times V \to W$ is just a function which is linear in the U and V arguments separately. A bilinear map $V \times \overline{V} \to W$ is often called *sesquilinear*.

[b] When $\Bbbk = \mathbb{R}$ the complex conjugate is of course unnecessary.

[c] When $\Bbbk = \mathbb{C}$ so that a priori we only have $(v, v) \in \mathbb{C}$, here the assertion $(v, v) \geq 0$ must be understood to mean that both $(v, v) \in \mathbb{R}$ and $(v, v) \geq 0$.

Equipping a vector space V with an inner product is commonly understood as giving V a rigid geometric structure: inner products allow us to measure lengths of vectors, and angles between pairs of vectors. For the future applications of inner products to vector bundles which lie in our future, we will instead be much more interested in a certain algebraic consequence of a choice of inner product on V, which we now describe.

Definition 2.31. A *complement* of a subspace U of a vector space V is another subspace $U' \subseteq V$ such that $U \cap U' = \{0\}$ and $\mathrm{span}(U \cup U') = V$.

Exercise 2.32. Let $U, U' \subseteq V$ be subspaces. We may define a linear map $f : U \oplus U' \to V$ by

$$u \oplus u' \mapsto u + u'.$$

Show that f is an isomorphism if and only if U' is a complement of U.

For this reason when U' is a complement of U it is sometimes said that V is the *internal direct sum* of U and U'.

Exercise 2.33. Let $U, U' \subseteq V$ be subspaces. A linear map $P : V \to V$ is called a *projection* if $P \circ P = P$. Show that U' is a complement of U if and only if there is a projection $P : V \to V$ with $\mathrm{im}\, P = U$ and $\ker P = U'$. Hint: use Exercise 2.32 to define a composite $P : V \to U \oplus U' \twoheadrightarrow U \hookrightarrow V$.

Proposition 2.34. *If* $(\cdot,\cdot) : V \times \overline{V} \to \Bbbk$ *is an inner product on a* $\Bbbk$*-vector space* V *then each subspace* $U \subseteq V$ *has a canonical complement.*

Exercise 2.35. Prove Proposition 2.34 by first showing that the set

$$U^{\perp} := \{v \in V : \forall u \in U, (v, u) = 0\}$$

is a subspace of V, and then showing that $U^{\perp}$ is a complement of U.

With respect to a fixed choice of inner product on V, the subspace $U^{\perp}$ is called the *orthogonal complement of* U.

Finally we observe the property (which will be especially useful when we get to vector bundles) that positive linear combinations of inner products remain inner products.

Exercise 2.36. Show that if $(\cdot,\cdot)_1$ and $(\cdot,\cdot)_2$ are both inner products on V (real or complex), then any positive linear combination $a(\cdot,\cdot)_1 + b(\cdot,\cdot)_2$ (i.e. with $a, b > 0$) is again an inner product on V.

The remainder of this section is completely optional and serves to give motivation for the definition of an inner product from entirely algebraic, rather than geometric, considerations. None of it is in any way necessary for our class on K-theory, but it does at least hopefully provide the keen reader with some interesting linear algebra exercises.

For simplicity, for the rest of this section suppose that V is a finite dimensional real vector space, so that an inner product on V is just the data of a bilinear form $V \times V \to \mathbb{R}$ (satisfying some conditions). Famously V and V^* are isomorphic, but not canonically so: let us at the outset fix an isomorphism $\Phi : V \to V^*$.

First notice that this choice of Φ at least produces something inner product-like: for $v_1, v_2 \in V$ we have that $\Phi(v_1) \in V^*$, which means that $\Phi(v_1)(v_2)$ is a real number. Thus we can define $(v_1, v_2) := \Phi(v_1)(v_2)$, but must keep in mind that this pairing need not be symmetric nor positive definite.

We would like the "generalized inner product" Φ which we are now considering to be as useful as an ordinary inner product. Perhaps the most useful property Φ could share with honest inner products would be for Φ to canonically produce complements of arbitrary subspaces $U \subseteq V$. For this purpose, as in Exercise 2.35 we define

$$U^\perp := \{v \in V : \forall u \in U, (v, u) = 0\}.$$

Because $(u, v) = \Phi(u)(v) \neq \Phi(v)(u) = (v, u)$ in general, we may also symmetrically define another orthogonal complement

$${}^\perp U := \{v \in V : \forall u \in U, (u, v) = 0\}.$$

on the other side. It turns out that, for the moment, $U^\perp$ will easier for us to use.

Definition 2.37. With respect to Φ, the subspaces $U^\perp$ and ${}^\perp U$ of V are the *left* and *right orthogonal complements of* U, respectively.

Exercise 2.38. Let V be finite dimensional and let $U \subseteq V$ be any subspace. Show that $\dim U^\perp = \dim {}^\perp U = \dim V - \dim U$. Hint: consider the linear map $\pi = \iota_U^* \circ \Phi : V \to U^*$ where ι_U^* was defined in Exercise 2.6. Show that $\ker \pi = U^\perp$ and $\operatorname{im} \pi = U^*$, and conclude by combining rank-nullity with Exercise 2.4. The case of the right orthogonal complement then follows formally from Proposition 2.45 below.

Exercise 2.39. Show that ${}^\perp(U^\perp) = U = ({}^\perp U)^\perp$. Hint: first show directly that $U \subseteq {}^\perp(U^\perp)$, and then conclude by Exercise 2.38. The other identity is analogous.

By definition, if $U^\perp$ is to be a complement of U in V then we first need $U \cap U^\perp = \{0\}$. In other words, we need that if $u \in U$ is such that $(u, u') = 0$ for all $u' \in U$, then actually $u = 0$. In the special case that $u' = u$ this in turn implies that if $\Phi(u)(u) = (u, u) = 0$, then $u = 0$. As the next exercise shows, it turns out that this last condition is equivalent to $U \cap U^\perp = \{0\}$.

Exercise 2.40. We have just seen that if $U \subseteq V$ is a subspace then the property "$(u,u) = 0$ implies $u = 0$" yields $U \cap U^{\perp} = \{0\}$. By considering the case of $U = \operatorname{span}(v)$ for fixed $v \in V$, show that the converse also holds.

Exercise 2.41. Show that the property "$(u,u) = 0$ implies $u = 0$" is also equivalent to $U \cap {}^{\perp}U = \{0\}$.

Thus from now on assume that Φ has the property:

$$\text{for all } v \in V, \text{ if } \Phi(v)(v) = 0 \text{ then } v = 0. \tag{2.4}$$

We will say that any isomorphism Φ which satisfies (2.4) is *definite*.

Exercise 2.42. Let $U \subset V$ be a subspace. We can then form the composite

$$\Phi_U : U \xhookrightarrow{\iota_U} V \xrightarrow{\Phi} V^* \xrightarrow{\iota_U^*} U^*$$

where the first map is the inclusion of U into V, the second map is just Φ, and the third map is the restriction map from Exercise 2.6.

Show that (2.4) implies that Φ_U is an isomorphism whenever Φ is an isomorphism.

Returning to the problem at hand, in order for $U^{\perp}$ to be a complement of U we also need that $V = \operatorname{span}(U \cup U^{\perp})$. In fact, we do not need to assume anything more to accomplish this.

Proposition 2.43. *Assuming that Φ is definite (i.e. obeys (2.4)), we have $V = \operatorname{span}(U \cup U^{\perp})$.*

Proof. By Exercise 2.33, it is sufficient to find a projection $P : V \to V$ with $\operatorname{im} P = U$ and $\ker P = U^{\perp}$. We construct such a P via the following abstract nonsense.

Lemma 2.44. *Let $\pi : V \to V'$ and $\sigma : V' \to V$ be linear maps such that $\pi \circ \sigma = \mathrm{id}_{V'}$. (A map σ with this property is called a section of π.) Then $P := \sigma \circ \pi$ is a projection $V \to V$ with $\ker P = \ker \pi$ and $\operatorname{im} P = \operatorname{im} \sigma$.*

Proof. We have $P \circ P = \sigma \circ (\pi \circ \sigma) \circ \pi = \sigma \circ \pi = P$ and so P is a projection. Now, since $\pi \circ \sigma = \mathrm{id}_{V'}$ we have that σ is injective and π is surjective. Therefore $\ker P = \ker \sigma \circ \pi = \ker \pi$ and $\operatorname{im} P = \operatorname{im} \sigma \circ \pi = \operatorname{im} \sigma$, as desired. □

Thus form the composites

$$\pi : V \xrightarrow{\Phi} V^* \xrightarrow{\iota_U^*} U^* \quad \text{and} \quad \sigma : U^* \xrightarrow{\Phi_U^{-1}} U \xhookrightarrow{\iota_U} V,$$

where $\iota_U : U \hookrightarrow V$ is the inclusion, $\iota_U^* : V^* \to U^*$ is the restriction map from Exercise 2.6, and Φ_U is the isomorphism from Exercise 2.42. Recall from Exercise 2.38

that $\ker \pi = U^\perp$ (if you like, because Φ is an isomorphism and $\ker \iota_U^* = \Phi(U^\perp)$). Moreover since Φ_U^{-1} is an isomorphism we have $\operatorname{im} \sigma = \operatorname{im} \iota_U = U$. Therefore Lemma 2.44 yields that $P = \sigma \circ \pi$ is a projection $V \to V$ with the desired properties. □

This all shows that $U^\perp$ is a complement of U in V, and therefore that definiteness of Φ implies that all subspaces of V are canonically complemented. The next proposition immediately implies that ${}^\perp U$ is also a complement of U.

Proposition 2.45. *If $\Phi : V \to V^*$ is a definite isomorphism (i.e. satisfying* (2.4)*), then that the linear map $\Psi : V \to V^*$ defined by*

$$\Psi(v)(v') := \Phi(v')(v)$$

is again a definite isomorphism. Moreover the left orthogonal complements with respect to Φ are right orthogonal complements with respect to Ψ, and vice versa.

Exercise 2.46. Prove Proposition 2.45.

The definiteness property (2.4) is reminiscent of the positive definiteness property $(v, v) \geq 0$ of inner products, but (2.4) drops the positivity requirement. The next exercise shows that definiteness is actually only a very minor weakening of positive definiteness.

Exercise 2.47. Show that if $\Phi : V \to V^*$ is a definite isomorphism (i.e. satisfying (2.4)), then the function $v \mapsto \Phi(v)(v)$ on V is either nonnegative or nonpositive. Hint: suppose that $\Phi(v)(v)$ and $\Phi(v')(v')$ have different signs and show that the function $f(t) = \Phi(tv + (1-t)v')(tv + (1-t)v')$ of the real variable t must pass through 0.

Finally, it is sometimes convenient to require that $U^\perp = {}^\perp U$ (since this need not occur in general). As we see in the next exercise, in this case we rediscover the notion of a symmetric pairing.

Exercise 2.48. Show that if $U^\perp = {}^\perp U$ for all subspaces $U \subseteq V$ then $\Phi(v)(v') = \Phi(v')(v)$ for all $v, v' \in V$. (The converse also holds.)

In summary, we have established the following.

Theorem 2.49. *Up to a sign, inner products on a real vector space V are in bijection with definite isomorphisms $V \to V^*$ which satisfy $U^\perp = {}^\perp U$ for all subspaces $U \subseteq V$.*

Exercise 2.50. Modify our arguments and constructions to handle the case of complex vector spaces. In particular, Proposition 2.45 and Exercise 2.48 need a small amount of fixing.

2.5 Universal properties

Just as we enumerated the universal properties of some common constructions in topology in Section 1, let us now briefly discuss the universal properties which our various linear algebraic constructions satisfy as well. We will also discuss the closely related examples of cases where linear maps between vector spaces induce linear maps between our constructions (e.g. that linear maps $U_1 \to V_1$ and $U_2 \to V_2$ canonically induce a linear map $U_1 \otimes U_2 \to V_1 \otimes V_2$). None of the details here will be necessary for our tutorial, but the definitions and proposition statements will likely be useful to keep in mind.

First recall our explicit construction of the tensor product in Subsection 2.2, and in particular note that we can trivially equip $U \otimes V$ with a bilinear map $\pi : U \times V \to U \otimes V$ defined by $(u, v) \mapsto u \otimes v$.

Proposition 2.51 (Universal property of tensor products)**.** *Let U and V both be vector spaces. For any other vector space W and bilinear map $B : U \times V \to W$ there exists a unique linear map $\widetilde{B} : U \otimes V \to W$ such that the diagram*

$$\begin{tikzcd} U \times V \arrow[r, "\pi"] \arrow[dr, "B"'] & U \otimes V \arrow[d, dashed, "\widetilde{B}"] \\ & W \end{tikzcd} \tag{2.5}$$

commutes. (The term commutes means that all possible composites of maps in the diagram which start and end at the same point give the same result.) In other words, such that $B = \widetilde{B} \circ \pi$.

Exercise 2.52. Prove Proposition 2.51.

This universal property lets us define the tensor product of linear maps $f_1 : U_1 \to V_1$ and $f_2 : U_2 \to V_2$ in the following way: observe that

$$(u_1, u_2) \mapsto f_1(u_1) \otimes f_2(u_2)$$

defines a bilinear map $B : U_1 \times U_2 \to V_1 \otimes V_2$. Then by the universal property of tensor products this bilinear map uniquely corresponds to a linear map $\widetilde{B} : U_1 \otimes U_2 \to V_1 \otimes V_2$.

Definition 2.53. The *tensor product* $f_1 \otimes f_2$ of $f_1 : U_1 \to V_1$ and $f_2 : U_2 \to V_2$ is the canonically induced map $U_1 \otimes U_2 \to V_1 \otimes V_2$ just defined.

Let us now briefly discuss the significance of universal properties in general. The universal property Proposition 2.51 actually completely characterizes the tensor product of two vector spaces in the following sense: suppose that T is any $\Bbbk$-vector space and $p : U \times V \to T$ is a bilinear map which together share the same universal property. That is, for each bilinear map $B : U \times V \to W$ there exists a unique linear map $\widehat{B} : T \to W$ such that $B = \widehat{B} \circ p$. Then in fact there is a canonical isomorphism between $U \times V$ and T, which must exist by the following general argument.

Apply the universal property of the tensor product for $U \otimes V$ to the bilinear map $p : U \times V \to T$ to produce a unique linear map $\widetilde{p} : U \otimes V \to T$ such that $p = \widetilde{p} \circ \pi$. Likewise apply the universal property of T to produce $\widehat{\pi} : T \to U \otimes V$. We claim that $\widetilde{p}$ and $\widehat{\pi}$ must be mutually inverse. To see this just note that $\widehat{\pi} \circ \widetilde{p}$ is a map $U \otimes V \to U \otimes V$ which satisfies

$$(\widehat{\pi} \circ \widetilde{p}) \circ \pi = \widehat{\pi} \circ p = \pi.$$

By the claimed uniqueness of such a map, since we also have $\mathrm{id}_{U\otimes V} \circ \pi = \pi$, we must have that $\mathrm{id}_{U\otimes V} = \widehat{\pi} \circ \widetilde{p}$. An analogous argument also shows that the composite $\widehat{\pi} \circ \widetilde{p}$ is the identity on T, as desired.

It follows from all of this that any “other” tensor product constructed by anyone else in any other way, which has merely also been show to satisfy the universal property of tensor products, must be canonically isomorphic to our own. Thus the upshot of the preceding is that once one knows that any $\Bbbk$-vector space T and bilinear map $p : U \times V \to p$ satisfying the universal property of tensor products for $U \otimes V$ exists at all, then one can safely forget about all of the details of how the pair (T, p) happened to be explicitly constructed. All useful properties of the tensor product necessarily follow from the universal property.

It’s not difficult to see how one generalizes the universal property of tensor products to kth tensor powers: one merely needs the concept of a *k-multilinear map* $M : V_1 \times \cdots V_k \to W$, i.e. which is linear in each of the k slots (we recover a bilinear form when $k = 2$). Let us now state the universal property precisely, along with the corresponding very similar universal properties of symmetric and exterior powers. When no confusion can arise, let us use π to denote all of the (distinct) natural maps from the k-fold product $V \times \cdots \times V$ into $V^{\otimes k}$, $S^k V$, and $\Lambda^k V$ respectively defined by mapping $(v_1, \ldots, v_k)$ to $v_1 \otimes \cdots \otimes v_k$, $v_1 \circledS \cdots \circledS v_k$, and $v_1 \wedge \cdots \wedge v_k$.

Proposition 2.54 (Universal property of tensor/symmetrc/exterior powers)**.** *Let V be a vector space. For any other vector space W and k-multilinear map $M : V \times \cdots \times V \to W$ there exists a unique linear map $\widetilde{M} : V^{\otimes k} \to W$ such that the diagram*

$$\begin{tikzcd} V \times \cdots \times V \arrow[r, "\pi"] \arrow[dr, "M"'] & V^{\otimes k} \arrow[d, dashed, "\widetilde{M}"] \\ & W \end{tikzcd}$$

commutes. The same result holds

- *if "k-multilinear" is replaced with "symmetric[8] k-multilinear" and $V^{\otimes k}$ is replaced with* $\mathrm{S}^k V$*, or*
- *if "k-multilinear" is replaced with "antisymmetric[9] k-multilinear" and $V^{\otimes k}$ is replaced with* $\Lambda^k V$*.*

Exercise 2.55. Prove Proposition 2.54.

In other words, symmetric k-multilinear maps $V \times \cdots \times V \to W$ are "the same" as linear maps $\mathrm{S}^k V \to W$, and likewise for antisymmetric k-multilinear maps and exterior powers.

As in the case of tensor products, these universal properties let us for example define the kth exterior power of a linear map $f : U \to V$: observe that the assignment

$$(u_1, \cdots, u_k) \mapsto f(u_1) \wedge \cdots \wedge f(u_k)$$

defines an antisymmetric bilinear map $M : U \times \cdots \times U \to \Lambda^k V$. Then by the universal property of exterior powers this bilinear map uniquely corresponds to a linear map $\widetilde{M} : \Lambda^k U \to \Lambda^k V$. There is of course a completely analogous construction for tensor and symmetric powers as well.

Definition 2.56. The *kth tensor* or *symmetric* or *exterior power* of $f : U \to V$, respectively denoted $f^{\otimes k}$ or $\mathrm{S}^k f$ or $\Lambda^k f$, is the canonically induced map $U^{\otimes k} \to V^{\otimes k}$ or $\mathrm{S}^k U \to \mathrm{S}^k V$ or $\Lambda^k U \to \Lambda^k V$ just defined.

In Exercise 2.22 we saw an alternative definition of the kth exterior power which worked when $\operatorname{char} \Bbbk \neq 2$: this other construction enjoys an analogous universal property to that of Proposition 2.54, with "k-multilinear" replaced with "alternating[10] k-multilinear" and $V^{\otimes k}$ replaced with our alternative construction of the kth exterior power.

Here are some related and more advanced exercises involving exterior powers.

Exercise 2.57. By analogy with Exercise 2.22 show that when $\operatorname{char} \Bbbk \neq 2$ we have that antisymmetric k-multilinear maps and alternating k-multilinear maps are the same. When $\operatorname{char} \Bbbk = 2$ then symmetric and antisymmetric k-multilinear maps become the same instead, while alternating k-multilinear maps are distinct in general.

[8] A *symmetric* k-multilinear map from a k-fold product $M : V \times \cdots \times V \to W$ is a map such that $M(v_1, \ldots, v_k) = M(v_{\sigma(1)}, \ldots, v_{\sigma(k)})$ whenever σ is a permutation of the numbers $1, \ldots, k$.

[9] An *antisymmetric* (or sometimes *skew*) k-multilinear map from a k-fold product $M : V \times \cdots \times V \to W$ is a map such that $M(v_1, \ldots, v_k) = \operatorname{sgn}(\sigma) M(v_{\sigma(1)}, \ldots, v_{\sigma(k)})$ whenever σ is a permutation of the numbers $1, \ldots, k$. Here $\operatorname{sgn}(\sigma)$ denotes the sign of the permutation σ.

[10] An *alternating* k-multilinear map from a k-fold product $M : V \times \cdots \times V \to W$ is a map such that $M(v_1, \ldots, v_k) = 0$ whenever $v_i = v_j$ for any $i \neq j$.

Exercise 2.58. Let V be a $\Bbbk$-vector space of dimension $n < \infty$.

1. Show that the set of antisymmetric k-multilinear maps $\mathrm{Alt}_k(V)$ on V is naturally a $\Bbbk$-vector space, which by Proposition 2.54 is canonically isomorphic to $(\Lambda^k V)^*$.

2. Show that for each $\psi_1 \wedge \cdots \wedge \psi_k \in \Lambda^k V^*$ we obtain an antisymmetric k-multilinear map by the definition

$$(v_1, \dots, v_k) \mapsto \sum_{\sigma \in S_k} \mathrm{sgn}(\sigma)\psi_1(v_{\sigma(1)}) \cdots \psi_k(v_{\sigma(k)}),$$

where S_k is the set of all permutations of $1, \dots, k$.

3. Moreover, show that the function $\Lambda^k V^* \to \mathrm{Alt}_k(V)$ defined by the previous part is a linear map. Conclude that by using the isomorphism $\mathrm{Alt}_k(V) \cong (\Lambda^k V)^*$ we have produced a canonical linear map $\Phi : \Lambda^k V^* \to (\Lambda^k V)^*$.

4. On the other hand, show that $(V^{\otimes k})^*$ and $(V^*)^{\otimes k}$ are canonically isomorphic.

5. Observe that each element of $(\Lambda^k V)^*$ gives rise to an element of $(V^{\otimes k})^*$ by precomposition with the canonical quotient map $\pi : V^{\otimes k} \to \Lambda^k V$. Thus show that we have a linear map

$$\Psi : (\Lambda^k V)^* \xrightarrow{\pi^*} (V^{\otimes k})^* \xrightarrow{\sim} (V^*)^{\otimes k} \xrightarrow{\pi} \Lambda^k V^*.$$

6. Show that $\Psi \circ \Phi = \frac{1}{k!}\,\mathrm{id}_{\Lambda^k V^*}$ and $\Phi \circ \Psi = \frac{1}{k!}\,\mathrm{id}_{(\Lambda^k V)^*}$. Conclude that when $\mathrm{char}\,\Bbbk > k$ then $\Lambda^k V^*$ and $(\Lambda^k V)^*$ are canonically isomorphic.

Exercise 2.59. Provide the analogous constructions for symmetric powers.

The direct sum of vector spaces can of course be characterized by a universal property too, but the correct universal property to write down depends on who you ask. (As it turns out, direct sums make sense in much greater generality, where there are a few different definitions which are in general not the same. However, in the special case of vector spaces all of these definitions become equivalent.) Here is one possible statement, which is very convenient but doesn't exactly have the same flavor as the universal properties above.

Proposition 2.60 (Universal property of direct sums)**.** *Let V_1 and V_2 all be vector spaces. There exist linear maps $\iota_i : V_i \hookrightarrow V_1 \oplus V_2$ and $\pi_i : V_1 \oplus V_2 \twoheadrightarrow V_i$ such that*

- *we have $\iota_1 \circ \pi_1 + \iota_2 \circ \pi_2 = \mathrm{id}_{V_1 \oplus V_2}$, and*

- *we have*

$$\pi_j \circ \iota_i = \begin{cases} \mathrm{id}_{V_i} & i = j \\ 0 & i \neq j \end{cases}.$$

Exercise 2.61. Define suitable maps ι_i and π_j and then prove Proposition 2.60.

We can use sums of the composites $\iota_i \circ \pi_j$ to define the direct sum of linear maps $f_1 : U_1 \to V_1$ and $f_2 : U_2 \to V_2$, but it is also not difficult to do this directly using our definition of the direct sum above.

Definition 2.62. The assignment

$$u_1 \oplus u_2 \mapsto f_1(u_1) \oplus f_2(u_2)$$

defines a linear map $f_1 \oplus f_2 : U_1 \oplus U_2 \to V_1 \oplus V_2$ called the *direct sum* of f_1 and f_2.

Finally, the evaluation $V^* \otimes V \to \Bbbk$ and coevaluation maps $\Bbbk \to V \otimes V^*$ of Exercise 2.14 and Exercise 2.15 respectively turn out to combine to completely characterize duals of finite-dimensional vector spaces. Let us now only sketch this imprecisely. To start, roughly speaking[a] (when V is finite-dimensional so that coev_V exists) we may form the composites[b]

$$V \longrightarrow \Bbbk \otimes V \xrightarrow{\mathrm{coev}_V \otimes \mathrm{id}_V} V \otimes V^* \otimes V \xrightarrow{\mathrm{id}_V \otimes \mathrm{ev}_V} V \otimes \Bbbk \longrightarrow V, \tag{2.6}$$

and

$$V^* \longrightarrow V^* \otimes \Bbbk \xrightarrow{\mathrm{id}_{V^*} \otimes \mathrm{coev}_V} V^* \otimes V \otimes V^* \xrightarrow{\mathrm{ev}_V \otimes \mathrm{id}_{V^*}} \Bbbk \otimes V^* \longrightarrow V^*. \tag{2.7}$$

It turns out that these composites are both always the (respective) identity map! In a more abstract setting, one actually says that *V has a left dual V^** if there exist maps ev_V and coev_V for which both of these composites are identity maps. One can then prove that if a left dual V^* of V exists then in fact V^* (along with the data of the maps ev_V and coev_V) is unique up to unique isomorphism. (There is actually a corresponding notion of *right duals*, but these two notions are the same for finite dimensional vector spaces.)

Of course, as stated the composites (2.6) and (2.7) do not actually make sense, because e.g. we are not able to actually form the tensor product $V \otimes V^* \otimes V$: we only know how to compute "binary tensor products", so we would have to choose to associate this as $(V \otimes V^*) \otimes V$ or $V \otimes (V^* \otimes V)$, and each choice results in a genuinely different object. These and other technical difficulties are resolved by introducing objects called *monoidal categories*, which abstractly axiomatize a collection of objects—along with the maps between them—equipped with a tensor product. A monoidal category where every object has duals (such as in the case of finite-dimensional vector spaces) is called *rigid*. When every object V is naturally[c] isomorphic to its double dual V^{**} then a rigid monoidal category is called *pivotal*.

Thus as a result of Exercise 2.7 we have seen that the category of finite dimensional vector spaces is a pivotal monoidal category.

Exercise 2.63. Show that there is a canonical isomorphism

$$\alpha_{U,V,W} : (U \otimes V) \otimes W \to U \otimes (V \otimes W),$$

called the *associator*. Also show that there is a canonical isomorphism $\beta_{U,V} : U \otimes V \to V \otimes U$ (called the *braiding*), for which the composite

$$U \otimes V \xrightarrow{\beta_{U,V}} V \otimes U \xrightarrow{\beta_{V,U}} U \otimes V$$

is always the identity. Such a braiding is called *symmetric*.

Exercise 2.64. Show that the isomorphism $\alpha_{U,V,W}$ is *natural in U* in the sense that for any linear map $f : U \to U'$ we have that the diagram

$$\begin{tikzcd}
(U \otimes V) \otimes W \arrow[r, "(f\otimes \mathrm{id}_V)\otimes \mathrm{id}_W"] \arrow[d, "\alpha_{U,V,W}"] & (U' \otimes V) \otimes W \arrow[d, "\alpha_{U',V,W}"] \\
U \otimes (V \otimes W) \arrow[r, "f\otimes(\mathrm{id}_V \otimes \mathrm{id}_W)"] & U' \otimes (V \otimes W)
\end{tikzcd}$$

commutes (and likewise for V and W). Also show the analogous property for $\beta_{U,V}$, i.e. that

$$\begin{tikzcd}
U \otimes V \arrow[r, "f\otimes \mathrm{id}_V"] \arrow[d, "\beta_{U,V}"] & U' \otimes V \arrow[d, "\beta_{U',V}"] \\
V \otimes U \arrow[r, "\mathrm{id}_V \otimes f"] & V \otimes U'
\end{tikzcd}$$

always commutes (and likewise for U).

Exercise 2.65. If V is finite-dimensional then we can form the composite

$$\Bbbk \xrightarrow{\mathrm{coev}_V} V \otimes V^* \xrightarrow{\beta_{V,V^*}} V^* \otimes V \xrightarrow{\mathrm{ev}_V} \Bbbk. \tag{2.8}$$

By Exercise 2.5 linear maps $\Bbbk \to \Bbbk$ canonically correspond to elements of $\Bbbk$: which element does the composite (2.8) correspond to? (Hint: this element is called the *categorical dimension* of V.)

Exercise 2.66. As an extension of the previous exercise, show that if $f : V \to V$ is a linear map then the composite

$$\Bbbk \xrightarrow{\mathrm{coev}_V} V \otimes V^* \xrightarrow{\beta_{V,V^*}} V^* \otimes V \xrightarrow{\mathrm{id}_{V^*} \otimes f} V^* \otimes V \xrightarrow{\mathrm{ev}_V} \Bbbk \tag{2.9}$$

corresponds to the trace of f. The composite (2.9) makes sense in greater generality (that we won't explore here) and is called the *categorical trace* of the map f.

Exercise 2.67. By inserting the isomorphism $\alpha_{U,V,W}$ from Exercise 2.63 (and its inverse) in necessary places, make the composites (2.6) and (2.7) actually make formal sense. (You'll have some choices when you do this, but all choices will give the same result.) Then prove that (2.6) and (2.7) are each equal to the identity.

Exercise 2.68. Existence of the evaluation and coevaluation maps actually allows us to define the dual f^* of a linear map $f : U \to V$ just by composition: roughly speaking, we form the composite

$$\begin{aligned} V^* \longrightarrow V^* \otimes \Bbbk &\xrightarrow{\mathrm{id}_{V^*} \otimes \mathrm{coev}_U} V^* \otimes U \otimes U^* \\ &\xrightarrow{\mathrm{id}_{V^*} \otimes f \otimes \mathrm{id}_{U^*}} V^* \otimes V \otimes U^* \\ &\xrightarrow{\mathrm{ev}_V \otimes \mathrm{id}_{U^*}} \Bbbk \otimes U^* \longrightarrow U^*. \end{aligned} \tag{2.10}$$

As in the previous exercise, insert the isomorphism $\alpha_{U,V,W}$ from Exercise 2.63 in necessary places so that (2.10) makes sense, and then show that the resulting composite is equal to f^* as defined in Exercise 2.6. Also note that there is a natural way to produce a map $V^* \to U^*$ in a similar way as (2.10) but by using coev_{U^*} and ev_{V^*} instead and also replacing f with its double dual $f^{**} : U^{**} \to V^{**}$: do this and show that you again recover the same map f^*.

This is just the beginning of the story: the isomorphism $\beta_{U,V}$ of Exercise 2.63 is a special case of the structure of a *braiding* on a monoidal category, which gives rise to the theory of braided and symmetric monoidal categories. The book to go for all of this (and much, much more) is Etingof–Gelaki–Nikshych–Ostrik [2].

[a] Up until this point it is unclear what a tensor product of $U \otimes V \otimes W$ of three vector spaces should actually mean: this is resolved in the coming exercises.

[b] Recall the isomorphism $V \to \Bbbk \otimes V$ and related maps from Exercise 2.13.

[c] Note that this word carries a precise meaning in category theory, and in fact in this case it is an open problem whether the word "naturally" is redundant here.

References

[1] Albrecht Dold. *Lectures on Algebraic Topology*. 20. Springer Berlin, 1980.

[2] Pavel Etingof, Shlomo Gelaki, Dmitri Nikshych, and Victor Ostrik. *Tensor Categories*. Vol. 205. American Mathematical Society, 2015. URL: `https://www-math.mit.edu/~etingof/egnobookfinal.pdf`.

[3] Allen Hatcher. *Algebraic Topology*. Springer Berlin, 2001. URL: `https://pi.math.cornell.edu/~hatcher/AT/AT.pdf`.

[4] James R Munkres. *Topology*. Vol. 2. Pearson Prentice Hall Upper Saddle River, 2000.