ZORN'S LEMMA AND SOME APPLICATIONS

KEITH CONRAD

1. Introduction

Zorn's lemma is a result in set theory that appears in proofs of some non-constructive existence theorems throughout mathematics. We will state Zorn's lemma below and use it in later sections to prove results in group theory, ring theory, linear algebra, and topology. In an appendix we give an application of Zorn's lemma to metric spaces. While Zorn's lemma has its origins in the early 20th century, it still retains a certain amount of "controversy" and we discuss why this is the case in the last section.

The statement of Zorn's lemma is not intuitive, and some of the terminology in it may be unfamiliar, but after reading through the explanation of Zorn's lemma and then the proofs that use it you should be more comfortable with how it can be applied.

Zorn's lemma: *Let S be a partially ordered set. If every totally ordered subset of S has an upper bound in S, then S contains a maximal element.*

To understand Zorn's lemma, we need to know four terms: partially ordered set, totally ordered subset, upper bound, and maximal element.

A *partial ordering* on a (nonempty) set S is a binary relation on S, denoted $\leq$, which satisfies the following properties:

- for all $s \in S$, $s \leq s$,
- if $s \leq s'$ and $s' \leq s$ then $s = s'$,
- if $s \leq s'$ and $s' \leq s''$ then $s \leq s''$.

When we fix a partial ordering $\leq$ on S, we refer to S (or, more precisely, to the pair $(S, \leq)$) as a partially ordered set.

It is important to notice that we do not assume all pairs of elements in S are comparable under $\leq$: for some s and s' we may have neither $s \leq s'$ nor $s' \leq s$. If all pairs of elements can be compared (that is, for all s and s' in S either $s \leq s'$ or $s' \leq s$) then we say S is *totally ordered* with respect to $\leq$.

Example 1.1. The usual ordering relation $\leq$ on $\mathbf{R}$ or on $\mathbf{Z}^+$ is a partial ordering of these sets. In fact it is a total ordering on either set. This ordering on $\mathbf{Z}^+$ is the basis for proofs by induction.

Example 1.2. On $\mathbf{Z}^+$, declare $a \leq b$ if $a \mid b$. This partial ordering on $\mathbf{Z}^+$ is different from the one in Example 1.1 and is called ordering by *divisibility*. It is one of the central relations in number theory. (Proofs about $\mathbf{Z}^+$ in number theory sometimes work not by induction, but by starting on primes, then extending to prime powers, and then extending to all positive integers using prime factorization. Such proofs view $\mathbf{Z}^+$ through the divisibility relation rather than through the usual ordering relation.) Unlike the ordering on $\mathbf{Z}^+$ in Example 1.1, $\mathbf{Z}^+$ is not totally ordered by divisibility: most pairs of integers are not comparable under the divisibility relation. For instance, 3 doesn't divide 5 and 5 doesn't divide 3. The subset $\{1, 2, 4, 8, 16, \dots\}$ of powers of 2 is totally ordered under divisibility.

Example 1.3. Let S be the set of all subgroups of a given group G. For $H, K \in S$ (that is, H and K are subgroups of G), declare $H \leq K$ if H is a subset of K. This is a partial ordering, called ordering by *inclusion*. It is usually not a total ordering: for most groups G there are subgroups H and K where $H \not\subset K$ and $K \not\subset H$.

Example 1.4. Let S be the set of linearly independent subsets of $\mathbf{R}^n$ (or another vector space). We can partially order S by inclusion: for two linearly independent subsets L and L' in $\mathbf{R}^n$, set $L \leq L'$ if $L \subset L'$.

Example 1.5. On $\mathbf{Z}^+$, declare $a \leq b$ if $b \mid a$. Here one positive integer is "larger" than another if it is a factor. This is called ordering by *reverse divisibility*.

Example 1.6. On the set of subgroups of a group G, declare subgroups H and K to satisfy $H \leq K$ if $K \subset H$. This is a partial ordering on the subgroups of G, called ordering by *reverse inclusion*.

In case you think ordering by reverse inclusion seems weird, let's take a look again at Example 1.2. There positive integers are ordered by divisibility, and nothing seems "backwards." But let's associate to each $a \in \mathbf{Z}^+$ the subgroup $a\mathbf{Z}$ of $\mathbf{Z}$. Every nonzero subgroup of $\mathbf{Z}$ has the form $a\mathbf{Z}$ for a unique positive integer a, $a\mathbf{Z} = b\mathbf{Z}$ if and only if $a = b$ (both a and b are positive), and $a \mid b$ if and only if $b\mathbf{Z} \subset a\mathbf{Z}$. For instance, $4 \mid 12$ and $12\mathbf{Z} \subset 4\mathbf{Z}$. Therefore the partial ordering on $\mathbf{Z}^+$ by divisibility (resp., reverse divisibility) is essentially the same as the partial ordering on nonzero subgroups of $\mathbf{Z}$ by reverse inclusion (resp., by inclusion). Partial ordering by reverse inclusion is used in the construction of completions of groups and rings.

Example 1.7. Let A and B be sets. Let S be the set of functions defined on some subset of A with values in B. The subset can vary with the function, but the codomain is always B. That is, S is the set of pairs (X, f) where $X \subset A$ and $f \colon X \to B$. Two elements (X, f) and (Y, g) in S are equal when $X = Y$ and $f(x) = g(x)$ for all $x \in X$.

We can partially order S by declaring $(X, f) \leq (Y, g)$ when $X \subset Y$ and $g|_X = f$. This means g is an extension of f to a larger subset of A. Let's check the second property of a partial ordering: if $(X, f) \leq (Y, g)$ and $(Y, g) \leq (X, f)$ then $X \subset Y$ and $Y \subset X$, so $X = Y$. Then the condition $g|_X = f$ means $g = f$ as functions on their common domain, with the same codomain B, so $(X, f) = (Y, g)$ in S.

Example 1.8. If S is a partially ordered set for the relation $\leq$ and $T \subset S$, then the relation $\leq$ provides a partial ordering on T. Thus T is a new partially ordered set under $\leq$. For instance, the partial ordering by inclusion on the subgroups of a group restricts to a partial ordering on the cyclic subgroups of a group.

In these examples, only Example 1.1 is totally ordered. This is typical: most naturally occurring partial orderings are not total orderings. However (and this is important) a partially ordered set can have many subsets that are totally ordered. As a dumb example, every one-element subset of a partially ordered set is totally ordered. A more interesting illustration is at the end of Example 1.2 with the powers of 2 inside $\mathbf{Z}^+$ under divisibility. As another example, if we partially order the subspaces of a vector space V by inclusion then a tower of subspaces

$$W_1 \subset W_2 \subset W_3 \subset \cdots$$

where each subspace is a proper subset of the next one is a totally ordered subset of V.

Here is a result about totally ordered subsets that will be useful at a few points later.

Lemma 1.9. *Let S be a partially ordered set. If $\{s_1, \dots, s_n\}$ is a finite totally ordered subset of S then there is an s_i such that $s_j \leq s_i$ for all $j = 1, \dots, n$.*

Proof. The s_i's are all comparable to each other; that's what being totally ordered means. Since we're dealing with a finite set of pairwise comparable elements, there will be one that is greater than or equal to them all in the partial ordering on S. The reader can formalize this with a proof by induction on n, or think about the bubble sort algorithm □

An *upper bound* on a subset T of a partially ordered set S is an $s \in S$ such that $t \leq s$ for all $t \in T$. When we say T has an upper bound in S, we do *not* assume the upper bound is in T itself; it is just in S.

Example 1.10. In $\mathbf{R}$ with its natural ordering, the subset $\mathbf{Z}$ has no upper bound while an upper bound on the subset of negative real numbers is 0 (or a positive real number). No upper bound on the negative real numbers is a negative real number.

Example 1.11. In the proper subgroups of $\mathbf{Z}$ ordered by inclusion, an upper bound on $\{4\mathbf{Z}, 6\mathbf{Z}, 8\mathbf{Z}\}$ is $2\mathbf{Z}$ since $4\mathbf{Z}$, $6\mathbf{Z}$, and $8\mathbf{Z}$ all consist entirely of even numbers. (Note $4\mathbf{Z} \subset 2\mathbf{Z}$, *not* $2\mathbf{Z} \subset 4\mathbf{Z}$.)

A *maximal* element m of a partially ordered set S is an element that is not below each element to which it is comparable: for all $s \in S$ to which m is comparable, $s \leq m$. Equivalently, m is maximal when the only $s \in S$ satisfying $m \leq s$ is $s = m$. This does *not* mean $s \leq m$ for all s in S since we don't insist that maximal elements are actually comparable to all elements of S.

Example 1.12. If we partially order linearly independent subsets of $\mathbf{R}^n$ by containment, then a maximal element is a linearly independent subset that is contained in no other linearly independent subset, and that is what a basis of $\mathbf{R}^n$ is: the maximal elements are the bases of $\mathbf{R}^n$.

Example 1.13. If we partially order $\mathbf{Z}^+$ by reverse divisibility (so $a \leq b$ means $b \mid a$), the number 1 is a maximal element. In fact 1 is the only maximal element. This is not a good example because 1 is comparable to everything in this relation, which is not a typical feature of maximal elements.

Example 1.14. Consider the positive integers *greater than* 1 with the reverse divisibility ordering: $a \leq b$ when $b \mid a$. The maximal elements here are the positive integers with no positive factor greater than 1 except themselves. These are the prime numbers, so the primes are the maximal elements for the reverse divisibility relation on $\{2, 3, 4, 5, 6, \dots\}$.

Equivalently, if we partially order the *proper* subgroups of $\mathbf{Z}$ by inclusion then the maximal elements are $p\mathbf{Z}$ for prime numbers p.

If $s \leq m$ for all $s \in S$ then we call m a *greatest element* of S. The special feature of a greatest element is being maximal *and* comparable to all of S. A greatest element of S, when it exists, is the only maximal element of S. A partially ordered set that has no greatest element could have many maximal elements. For example, the subgroups of $\mathbf{Z}$ partially ordered by containment have $\mathbf{Z}$ as its greatest (and only maximal) element, while the *proper* subgroups of $\mathbf{Z}$ partially ordered by inclusion have no greatest element and have many maximal elements: subgroups $p\mathbf{Z}$ for prime p. We will not be concerned with greatest elements and it's best to forget about them here.

We now return to the statement of Zorn's lemma:

> If every totally ordered subset of a partially ordered set S has an upper bound in S, then S contains a maximal element.

All the terms being used here have now been defined.[1] Of course this doesn't mean the statement should be clearer!

Zorn's lemma is very nonconstructive: when it can be applied, it provides no mechanism to find a maximal element whose existence it asserts and it says nothing about how many maximal elements there are. Usually, as in Example 1.14, there are many maximal elements.

In a partially ordered set S we can speak about minimal elements just as much as maximal elements: $m \in S$ is called *minimal* if $m \leq s$ for all $s \in S$ to which m is comparable. Zorn's lemma can be stated in terms of minimal elements: if every totally ordered subset of a partially ordered set S has a lower bound in S then S has a minimal element. There really is no need to use this formulation in practice, since by reversing the meaning of the partial ordering (that is, using the reverse ordering) lower bounds become upper bounds and minimal elements become maximal elements. Analogous to greatest elements are least elements: $m \in S$ is a *least element* if $m \leq s$ for all $s \in S$, so m is both minimal and comparable to all elements of S. In $\mathbf{Z}^+$ ordered in the standard way, each nonempty subset has a least element but not necessarily a greatest element.

Zorn's lemma is not intuitive, but it is logically equivalent to more intuitively plausible statements in set theory like the Axiom of Choice, which says every Cartesian product of nonempty sets is nonempty: if X_i $(i \in I)$ are nonempty sets then $\prod_{i \in I} X_i \neq \emptyset$. In the set theory appendix to [21], Zorn's lemma is derived from the Axiom of Choice. The equivalence between Zorn's lemma and the Axiom of Choice is proved in an appendix to [25]. The reason for calling Zorn's lemma a lemma rather than an axiom is purely historical. Zorn's lemma is also equivalent to the Well-Ordering theorem, which says every nonempty set has a well-ordering: that is a total ordering on the set such that every nonempty subset has a least element. Zorn's lemma was introduced by Max Zorn in 1935 [32] to shorten proofs in algebra that previously had used the Axiom of Choice or the Well-Ordering theorem.

We will discuss uses of Zorn's lemma mostly in algebra, but it shows up in many other areas. For instance, the most important result in functional analysis is the Hahn-Banach theorem, whose proof uses Zorn's lemma. Another result from functional analysis, the Krein-Milman theorem, is proved using Zorn's lemma. (The Krein-Milman theorem is an example where Zorn's lemma proves the existence of something that is more naturally a minimal element rather than a maximal element.) In topology, the most important theorem about compact spaces is Tychonoff's theorem, and its proof uses Zorn's lemma.

When dealing with objects that have a built-in finiteness condition (such as finite-dimensional vector spaces or finite products of spaces $X_1 \times \cdots \times X_n$), Zorn's lemma can be avoided by using ordinary induction in a suitable way (*e.g.*, inducting on the dimension of a vector space). The essential uses of Zorn's lemma are for truly infinite objects, where one has to make infinitely many choices at once in a rather extreme way. Tim Gowers[2] gave a nice description of when you can anticipate a proof will use Zorn's lemma:

> Typically, one is trying to build a structure of some kind [...]. The natural way to do it appears to be to build the structure up in stages, but there are too many stages for this to work straightforwardly. However, once one has

[1] The hypotheses refer to *all* totally ordered subsets of S, and a totally ordered subset might be uncountable. Therefore it is a bad idea to write about "totally ordered sequences," since the label "sequence" is often understood to refer to a countably indexed set. Just use the label "totally ordered subset."

[2] See the end of https://gowers.wordpress.com/2008/08/12/how-to-use-zorns-lemma/.

> an idea of what a stage is and what the building-up process is, one can wheel out Zorn's lemma to finish the job. The partially ordered set will consist of all objects that might conceivably be stages in the construction, and one of these objects will be smaller than another if it might conceivably come before the other in the building-up process. If the resulting partial order satisfies the chain condition and if a maximal element must be a structure of the kind one is trying to build, then the proof is complete.

You will see this idea in action in the proofs coming up.

2. Applications to group theory

There are two common reasons that Zorn's lemma is used: to find a subset that is as big as possible subject to some conditions or to show a function defined on a subset can be extended to the whole set while preserving certain technical conditions. We will illustrate both ideas in this section within the setting of group theory.

A subgroup M of a nontrivial group G is called a *maximal subgroup* if M is a proper subgroup of G and there is no subgroup strictly contained between M and G. This means M is maximal for containment among the proper subgroups of G. (You may think a better term would be "maximal proper subgroup" instead of "maximal subgroup," but the term "maximal subgroup" is standard.)

Example 2.1. Let's show the maximal subgroups of $\mathbf{Z}$ are the subgroups $p\mathbf{Z}$ for prime p. A proper subgroup of $\mathbf{Z}$ is $m\mathbf{Z}$ where $m \neq \pm 1$, and $m\mathbf{Z} \subset p\mathbf{Z}$ where p is a prime factor of m. Therefore a maximal subgroup of $\mathbf{Z}$ must be $p\mathbf{Z}$ for a prime p. Conversely, $p\mathbf{Z}$ for each prime p is a maximal subgroup of $\mathbf{Z}$ since if $a\mathbf{Z}$ is a subgroup of $\mathbf{Z}$ such that $p\mathbf{Z} \subset a\mathbf{Z} \subset \mathbf{Z}$ then $a \mid p$, so $a = \pm 1$ or $\pm p$. Thus $a\mathbf{Z}$ equals $\mathbf{Z}$ or $p\mathbf{Z}$.

Example 2.2. The additive group $\mathbf{Q}$ has *no* maximal subgroups. For every proper subgroup H of $\mathbf{Q}$, we'll show there is a subgroup strictly between H and $\mathbf{Q}$. First, $[\mathbf{Q} : H] = \infty$: if the index were finite, say $n > 1$, then the quotient group $\mathbf{Q}/H$ would have order n, so $nr \equiv 0 \bmod H$ for all $r \in \mathbf{Q}$. That means $n\mathbf{Q} \subset H$. Since $n\mathbf{Q} = \mathbf{Q}$, we get $\mathbf{Q} \subset H$ and thus $H = \mathbf{Q}$, which contradicts $[\mathbf{Q} : H] > 1$. Pick $r \in \mathbf{Q} - H$, so $r \neq 0$. The subgroup $H + \mathbf{Z}r$ of $\mathbf{Q}$ strictly contains H and we'll show $[H + \mathbf{Z}\langle r\rangle : H]$ is finite. Representatives for the quotient group $(H + \mathbf{Z}\langle r\rangle)/H$ can be chosen from integral multiples of r, and r in $(H + \mathbf{Z}\langle r\rangle)/H$ has finite order: write $r = a/b$ for $a, b \in \mathbf{Z} - \{0\}$ and pick a nonzero $h = c/d$ in H, where $c, d \in \mathbf{Z} - \{0\}$. Then $bcr = ac = adh \in H$, so r in $(H + \mathbf{Z}\langle r\rangle)/H$ has order at most $|bc|$. Thus $(H + \mathbf{Z}r)/H$ is finite while $\mathbf{Q}/H$ is infinite, so $H \subsetneqq H + \mathbf{Z}\langle r\rangle \subsetneqq \mathbf{Q}$.

The following theorem gives a condition under which a group has maximal subgroups.

Theorem 2.3. *Let G be a nontrivial finitely generated group. Every proper subgroup of G is contained in a maximal subgroup of G.*

This theorem does not apply to $\mathbf{Q}$ since it is not finitely generated.

Proof. Let $g_1, \ldots, g_n$ be a finite (nonempty) generating set for G, so $G = \langle g_1, \ldots, g_n\rangle$. For a proper subgroup $\widetilde{G}$ of G and $0 \leq j \leq n$, set $G_j = \langle \widetilde{G}, g_1, \ldots, g_j\rangle$, so

$$\widetilde{G} = G_0 \subset G_1 \subset \cdots \subset G_n = G.$$

Let ℓ be the largest integer from 0 to $n-1$ such that $G_\ell \neq G$. Then $G_{\ell+1} = \langle G_\ell, g_{\ell+1}\rangle = G$, so $g_{\ell+1} \notin G_\ell$.

Define a partially ordered set. Let

$$S = \{H : H \text{ is a subgroup of } G, G_\ell \subset H, g_{\ell+1} \notin H\},$$

where subgroups in S are ordered by inclusion. The set S is nonempty since $G_\ell \in S$. We'll apply Zorn's lemma to S to get a maximal subgroup of G that contains $\widetilde{G}$.

Totally ordered subsets of S have an upper bound in S. Let $\{H_i\}_{i\in I}$ be a totally ordered subset of S. We will show the union $H = \bigcup_{i\in I} H_i$ is an upper bound in S on every H_i. Even though a union of subgroups need not be a subgroup in general (think about $2\mathbf{Z} \cup 3\mathbf{Z}$ inside $\mathbf{Z}$), the fact that the subgroups H_i are totally ordered will make H a subgroup.

If h and h' are in H then $h \in H_i$ and $h' \in H_{i'}$ for some subgroups H_i and $H_{i'}$. Since we are working with a totally ordered set of subgroups in S, $H_i \subset H_{i'}$ or $H_{i'} \subset H_i$. Without loss of generality, $H_i \subset H_{i'}$. Therefore h and h' are in $H_{i'}$, so $hh' \in H_{i'} \subset H$. Also $h^{-1} \in H_i \subset H$. Therefore H is a subgroup of G. We have $H \in S$ since (i) $g_{\ell+1} \notin H$ and (ii) $G_\ell \subset H$. To prove (i), suppose $g_{\ell+1} \in H$. Then $g_{\ell+1} \in H_i$ for some i, but that contradicts the fact that H_i belongs to S. So $g_{\ell+1} \notin H$. To prove (ii), since G_ℓ is contained in each H_i (or even just one of them), G_ℓ is contained in their union, so $G_\ell \subset H$. Thus $H \in S$.

Since H contains each H_i and $H \in S$, H is an upper bound in S on $\{H_i\}_{i\in I}$.

Use Zorn's lemma. By Zorn's lemma, S has a maximal element M: M is a subgroup of G with $G_\ell \subset M$ and $g_{\ell+1} \notin M$, and M is maximal for containment among all subgroups of G *with those two properties*.

M is a maximal subgroup of G and contains $\widetilde{G}$. Since $\widetilde{G} \subset G_\ell \subset M$, M contains $\widetilde{G}$. For M to be a maximal subgroup of G means M is not contained in another proper subgroup of G, with no further conditions: if $M \subset H \subset G$ for a subgroup $H \neq M$ then $H = G$.

From $G_\ell \subset M \subset H$, we have $G_\ell \subset H$. To show $H = G$, suppose $g_{\ell+1} \notin H$. Then $H \in S$ (see the definition of S), and this contradicts the maximality of M in S. Therefore $g_{\ell+1} \in H$, so $\langle M, g_{\ell+1}\rangle \subset H$. Since M contains G_ℓ, $\langle M, g_{\ell+1}\rangle$ contains $\langle G_\ell, g_{\ell+1}\rangle = G_{\ell+1} = G$, so $G \subset H$. Thus $H = G$. □

The proof of Theorem 2.3 exhibits a standard *disconnect* between upper bounds on totally ordered subsets and maximal elements in the whole set. Consider proper subgroups of $\mathbf{Z}$ under inclusion. The maximal subgroups are $p\mathbf{Z}$ for prime numbers p. The subgroups $\{6\mathbf{Z}, 12\mathbf{Z}, 24\mathbf{Z}\}$ are totally ordered under inclusion, and in the proof of Theorem 2.3 the upper bound created on this subset is the union $6\mathbf{Z} \cup 12\mathbf{Z} \cup 24\mathbf{Z} = 6\mathbf{Z}$. (A finite totally ordered subset of a partially ordered set always has one of its members as an upper bound on the subset, by Lemma 1.9.) This upper bound is *not* a maximal subgroup of $\mathbf{Z}$. So the task of checking the hypotheses of Zorn's lemma are satisfied is a completely separate matter from applying Zorn's lemma: an upper bound on a totally ordered subset does not have to be a maximal element of the whole set. Remember that!

Our next application of Zorn's lemma is to extending the domain of a homomorphism. An abelian group D is called *divisible* if the function $x \mapsto nx$ is surjective for every $n \geq 1$. (We write the group operation additively.) That is, for each $d \in D$ and $n \geq 1$ there is an x such that $d = nx$, so we can "divide" d by n (but x need not be unique). For example, $\mathbf{R}/\mathbf{Z}$ and $\mathbf{Q}$ are divisible groups while $\mathbf{Z}$ is not. Among multiplicative groups, $\mathbf{C}^\times$ and its subgroup S^1 are divisible. (The function $x \mapsto e^{2\pi i x}$ sets up an isomorphism $\mathbf{R}/\mathbf{Z} \cong S^1$.)

Theorem 2.4. *Let D be a divisible group. If A is an abelian group and $B \subset A$ is a subgroup, each homomorphism $f\colon B \to D$ can be extended to a homomorphism $\widetilde{f}\colon A \to D$.*

Proof. Pick $a \in A$ with $a \notin B$. Then the subgroup $\langle B, a\rangle = B + \mathbf{Z}a$ spanned by B and a contains B. As a warm-up we will show how to extend f to this larger subgroup of A. (See the diagram below.) Then we will bring in Zorn's lemma.

$$\begin{tikzcd}
A \arrow[d, dash] & \\
\langle B, a\rangle \arrow[d, dash, "\neq"'] \arrow[dr, dotted] & \\
B \arrow[r, "f"'] & D.
\end{tikzcd}$$

Consider how $\langle a\rangle$ can meet B. The set $\{k \in \mathbf{Z} : ka \in B\}$ is a subgroup of $\mathbf{Z}$, so it is 0 or $n\mathbf{Z}$ for some $n \geq 1$.

Case 1. $\{k \in \mathbf{Z} : ka \in B\} = 0$. Then each element of $\langle B, a\rangle$ is $b + ka$ for unique $b \in B$ and $k \in \mathbf{Z}$ (why?). Define $f' \colon \langle B, a\rangle \to D$ by $f'(b + ka) = f(b)$, so $f'|_B = f$. The mapping f' is a homomorphism since

$$\begin{aligned}
f'((b_1 + k_1a) + (b_2 + k_2a)) &= f'((b_1 + b_2) + (k_1 + k_2)a) \\
&= f(b_1) + f(b_2) \\
&= f'(b_1 + k_1a) + f'(b_2 + k_2a).
\end{aligned}$$

Case 2. $\{k \in \mathbf{Z} : ka \in B\} = n\mathbf{Z}$ for some $n \geq 1$. Then a positive multiple of a lies in B, and na is that multiple with n minimal. The function f makes sense at na, but not at a. If we can extend f to a homomorphism $f' \colon \langle B, a\rangle \to D$ then $f'(a)$ must satisfy the relation $nf'(a) = f'(na) = f(na)$. Here $f(na)$ is already defined while $f'(a)$ is not. To define $f'(a)$ we need to find an $x \in D$ such that $nx = f(na)$ and then we set $f'(a)$ to be that x. Because D is divisible, there is an $x \in D$ such that $nx = f'(a)$. Define $f'(a) = x$, and more generally define

$$f'(b + ka) = f(b) + kx.$$

Is this well-defined? Suppose

$$b + ka = b' + k'a. \tag{2.1}$$

Then $(k - k')a = b' - b \in B$, so $k - k' \in n\mathbf{Z}$ by the definition of n. Write $k = k' + n\ell$ for some $\ell \in \mathbf{Z}$, so

$$\begin{aligned}
f(b) + kx &= f(b) + (k' + n\ell)x \\
&= f(b) + k'x + \ell(nx) \\
&= f(b) + k'x + \ell f(na) \\
&= f(b + \ell na) + k'x.
\end{aligned}$$

Since $b + \ell na = b + (k - k')a = b + ka - k'a$, and that's b' by (2.1). Thus

$$f(b) + kx = f(b') + k'x,$$

so $f' \colon \langle B, a\rangle \to D$ is well-defined. It is left to the reader to show it is a homomorphism.

Now we show Zorn's lemma can be applied.

Define a partially ordered set. Let S be the set of pairs (C, g) where C is a subgroup between B and A and $g \colon C \to D$ is a homomorphism that extends f (that is, $g|_B = f$). The picture is as follows.

$$\begin{tikzcd}
A \arrow[d, dash] & \\
C \arrow[d, dash] \arrow[dr, "g"] & \\
B \arrow[r, "f"'] & D.
\end{tikzcd}$$

The set S is nonempty since $(B, f) \in S$. Partially order S by declaring $(C, g) \leq (C', g')$ if $C \subset C'$ and $g'|_C = g$. That is, g' extends g to the larger intermediate subgroup C'.

$$\begin{tikzcd}
A \arrow[d, dash] & \\
C' \arrow[d, dash] \arrow[ddr, "g'"] & \\
C \arrow[d, dash] \arrow[dr, "g"] & \\
B \arrow[r, "f"'] & D.
\end{tikzcd}$$

Totally ordered subsets of S have an upper bound in S. If $\{(C_i, g_i)\}_{i \in I}$ is a totally ordered subset of S, then it has an upper bound in S: use $C = \bigcup_{i \in I} C_i$ as the subgroup (it really is a subgroup, using an argument similar to $\bigcup_{i \in I} H_i$ being a subgroup in the proof of Theorem 2.3) and let $g\colon C \to D$ by $g(x) = g_i(x)$ if $x \in C_i$. Is this well-defined? Well, supposing x is in C_i and C_j, we need to know $g_i(x) = g_j(x)$. Either $(C_i, g_i) \leq (C_j, g_j)$ or $(C_j, g_j) \leq (C_i, g_i)$ since the (C_i, g_i)'s are totally ordered in S. If $(C_i, g_i) \leq (C_j, g_j)$ then $C_i \subset C_j$ and $g_j|_{C_i} = g_i$, so $g_j(x) = g_i(x)$. That $g_i(x) = g_j(x)$ if $(C_j, g_j) \leq (C_i, g_i)$ is proved in the same way. Because each pair of elements in C lies in a common C_i by Lemma 1.9, g is a homomorphism because each g_i is a homomorphism. Since $g_i|_B = f$ for all i, $g|_B = f$. Thus $(C, g) \in S$ and (C, g) is an upper bound on $\{(C_i, g_i)\}_{i \in I}$.

Use Zorn's lemma. The hypotheses of Zorn's lemma on S have been checked, so S has a maximal element (M, h). That is, M is a group between B and A, $h\colon M \to D$ is a homomorphism and there is no extension of h to a homomorphism out of a larger subgroup of A than M.

$M = A$, so h is defined on all of A. We will show this by *contradiction*. If $M \neq A$ then there is some $a \in A$ with $a \notin M$. By the argument used at the start of this proof, with B and f there replaced by M and h here, h can be extended to a homomorphism $h'\colon \langle M, a \rangle \to D$.

$$\begin{tikzcd}
A \arrow[d, dash] & \\
\langle M, a \rangle \arrow[d, dash, "\neq"'] \arrow[dr, dotted, "h'"] & \\
M \arrow[r, "h"'] & D.
\end{tikzcd}$$

Thus $(M, h) \leq (\langle M, a\rangle, h')$ and $(M, h) \neq (\langle M, a\rangle, h')$, which contradicts the maximality of (M, h). Hence $M = A$. □

Corollary 2.5. *Let D be a divisible group. If D is a subgroup of an abelian group A then D is a direct factor in A: there is a direct product decomposition $A = D \times H$ (internal direct product) for some subgroup H of A.*

Proof. The identity function $f\colon D \to D$ is a homomorphism and D is a subgroup of A, so f can be extended to a homomorphism $\widetilde{f}\colon A \to D$ by Zorn's lemma, thanks to Theorem 2.4. We will show $A = D \times \ker(\widetilde{f})$.

For $a \in A$, set $d = \widetilde{f}(a)$. Since $\widetilde{f}(d) = f(d) = d$, $\widetilde{f}(a) = \widetilde{f}(d)$. Thus $h := a - d$ is in $\ker(\widetilde{f})$ and $a = d + h$. We have shown $A = D + \ker(\widetilde{f})$. This decomposition is direct since $D \cap \ker(\widetilde{f}) = \{0\}$: if $d \in D \cap \ker(\widetilde{f})$ then $0 = \widetilde{f}(d) = f(d) = d$. □

Example 2.6. Since S^1 is a divisible subgroup of $\mathbf{C}^\times$, $\mathbf{C}^\times = S^1 \times H$ for a subgroup H of $\mathbf{C}^\times$. Indeed, we can use $H = \mathbf{R}_{>0}$: $\mathbf{C}^\times = S^1 \times \mathbf{R}_{>0}$ by writing nonzero z in $\mathbf{C}$ as $e^{i\theta}r$ for $r = |z|$ and $e^{i\theta} = z/|z| \in S^1$. This is a situation where we don't really need Corollary 2.5 to get a direct product decomposition of a group with the subgroup S^1 as one factor.

Example 2.7. Let μ_∞ be the group of all roots of unity in $\mathbf{C}^\times$. (A standard notation for the nth roots of unity in $\mathbf{C}^\times$ is μ_n.) Since μ_∞ is a divisible subgroup of S^1, $S^1 = \mu_\infty \times H$ for a subgroup H of S^1. In contrast to the previous example, the direct product decomposition here is completely abstract and we only know it exists from the proof of Corollary 2.5, which uses Zorn's lemma. Specifically, in that proof there is no known concrete way to write down a group homomorphism $S^1 \to \mu_\infty$ that is the identity on μ_∞.

Example 2.8. Since $\mathbf{Q}$ is a divisible subgroup of $\mathbf{R}$, $\mathbf{R} = \mathbf{Q} \times H$ for a subgroup H of $\mathbf{R}$. (This is an additive decomposition of $\mathbf{R}$, not a multiplicative one, so writing $\mathbf{R} = \mathbf{Q} \oplus H$ might be better.) No such decomposition of $\mathbf{R}$ is known where H can be described explicitly.

Remark 2.9. The application of Zorn's lemma in Theorem 2.4 generalizes from abelian groups ($\mathbf{Z}$-modules) to modules over a commutative ring R, and is called Baer's characterization of injective R-modules. (An R-module is injective when it is a direct summand of each module it can be embedded in.) See [10, p. 396] or [25, p. 483]. Divisible abelian groups are examples of injective $\mathbf{Z}$-modules.

Let's use Zorn's lemma to do something crazy: show there is a "maximal" group. On the set of all groups, define the partial ordering by inclusion. This is a partial ordering. If $\{G_i\}$ is a totally ordered set of groups, let G be the union of the G_i's. Every pair of elements in G is in a common G_i, so it is easy to define the group law in G (use the group law in G_i), check it is well-defined (independent of the G_i containing the two elements), and G contains all G_i's. So now it seems, by Zorn's lemma, that we should have a maximal group: a group that is not contained in some larger group. This is absurd, since for each group G we can create $G \times \mathbf{Z}$ and literally (if you wish) replace the elements of $G \times \{0\}$ with the elements of G to make G a genuine subset of $G \times \mathbf{Z}$. Then G is properly contained in another group and we have contradicted maximality. What is the error?

The problem here is right at the start when we defined our partially ordered set as the "set of all groups." This is the kind of set-theoretic looseness that leads to paradoxes (set of all sets, *etc.*). In fact, the group-theoretic aspect of the construction was irrelevant for the contradiction: if we just worked with sets and containment, the same argument goes through to show there is a set contained in no other sets, which is false and contains the same error as above: there is no "set of all sets" that one can partially order to make Zorn's lemma apply. Looking back now, perhaps with some suspicion, at the previous two proofs where we created an upper bound on a totally ordered subset of a partially ordered set relative to some kind of inclusion relation, the objects that we formed the union of were subsets of a larger fixed set (subgroups of a group). In that context the union really makes sense. It doesn't make sense to take arbitrary unions of arbitrary sets that *a priori* don't live in some common set. As an extreme case, we can't take the union of "all sets" to find one set containing all others.

Here are a few exercises about Zorn's lemma using ideas from this section.

Exercise. At the start of this section we saw $\mathbf{Q}$ has no maximal proper subgroup. By imposing a constraint on the kind of subgroup, we can obtain maximal examples.

a) Use Zorn's lemma to show the set of subgroups of $\mathbf{Q}$ that *don't contain* 1 has maximal elements: there is a subgroup M of $\mathbf{Q}$ such that (i) $1 \notin M$ and (ii) if $M \subset H \subset \mathbf{Q}$ and $1 \notin H$ then $H = M$. (Start by showing the set of such subgroups is not empty!)

b) Let M be a subgroup of $\mathbf{Q}$ maximal for the property of not containing 1. Show $M \cap \mathbf{Z} \neq \{0\}$ (hint: if $M \cap \mathbf{Z} = \{0\}$ then consider $M + n\mathbf{Z}$ where $n \geq 2$) and then show $M \cap \mathbf{Z} = q\mathbf{Z}$ for a prime q.

c) Show the numerator of each fraction in M is divisible by q from (b), so $M \subset H_q := \{a/b : a, b \in \mathbf{Z}, q \mid a, q \nmid b\}$. Elements of H_q are the fractions that "make sense" mod q.

d) For every prime p, define $H_p := \{a/b : a, b \in \mathbf{Z}, p \mid a, p \nmid b\}$. Prove H_p is a subgroup of $\mathbf{Q}$, $H_p \cap \mathbf{Z} = p\mathbf{Z}$, and H_p is maximal among subgroups of $\mathbf{Q}$ not containing 1. Therefore by (c), the maximal subgroups in (a) are the subgroups H_p.

e) Give an example of a proper subgroup of $\mathbf{Q}$ containing H_p. (It must contain 1.)

Exercise. Let A be an abelian group, written additively, such that all elements have order dividing a fixed positive integer m: $ma = \{0\}$ for all $a \in A$. An example is an arbitrary direct product $A = \prod_{j \in J} \mathbf{Z}/(m)$, but examples can be far more complicated.

a) Suppose B is a subgroup of A and there is a homomorphism $f \colon B \to \mathbf{Z}/(m)$. For each $a \in A - B$, show f can be extended to a homomorphism $\langle B, a \rangle \to \mathbf{Z}/(m)$.

b) Use Zorn's lemma to show for each subgroup B of A and homomorphism $f \colon B \to \mathbf{Z}/(m)$ that f can be extended to a homomorphism $A \to \mathbf{Z}/(m)$.

c) If there is an $a \in A$ with order m, then use (b) to show $A = \langle a \rangle \times A'$ for some subgroup A' of A.

Exercise. Let K and F be fields. There need not be a homomorphism $K \to F$ (*e.g.*, $K = \mathbf{Q}$ and $F = \mathbf{F}_2$). But assume some subring of K admits a ring homomorphism to F (*e.g.*, if K has characteristic 0 then $\mathbf{Z} \subset K$ and there is certainly a ring homomorphism $\mathbf{Z} \to F$, while if K and F both have characteristic $p > 0$ then $\mathbf{Z}/(p)$ is a subfield of K and F so there is an inclusion homomorphism $\mathbf{Z}/(p) \hookrightarrow F$).

a) Let S be the set of pairs (A, f) where A is a subring of K and $f \colon A \to F$ is a ring homomorphism, so $S \neq \emptyset$ by hypothesis. Partially order S by $(A, f) \leq (B, g)$ if $A \subset B$ and $g|_A = f$. Show with Zorn's lemma that S contains a maximal pair, which amounts to a subring $A \subset K$ that admits a ring homomorphism $f \colon A \to F$ that can't be extended to a ring homomorphism out of a larger subring of K.

b) When $K = \mathbf{Q}$ and $F = \mathbf{Z}/(2)$, show the ring of fractions $\{m/n : m \in \mathbf{Z}, n \in \mathbf{Z} - \{0\}, n \text{ is odd}\}$ admits a homomorphism to $\mathbf{Z}/(2)$ and is the ring part of a maximal pair

(A, f) in S. (Note $\mathbf{Q}$ itself is not part of a maximal pair since there is no ring homomorphism from $\mathbf{Q}$ to $\mathbf{Z}/(2)$.) What is the kernel of this homomorphism?

c) If (A, f) is a maximal pair in S, show every element of A not in $\ker f$ is a unit in A, and therefore $\ker f$ is the only maximal ideal in A. (A ring with only one maximal ideal is called a *local ring*. Trying to find a maximal subring of one field that admits a homomorphism to another field gives rise, using Zorn's lemma, to examples of local rings.)

3. Applications to ideals

The ideals in a commutative ring can be partially ordered by inclusion. The whole ring, which is the unit ideal (1), is obviously maximal with respect to this ordering. But this is boring and useless. Proper ideals that are maximal for inclusion among the proper ideals are called the maximal ideals in the ring. (That is, a maximal ideal is understood to mean a maximal proper ideal.) Let's prove they always exist.

Theorem 3.1 (Krull). *Every nonzero commutative ring contains a maximal ideal.*

Proof. Let S be the set of proper ideals in a commutative ring $R \neq 0$. Since the zero ideal (0) is a proper ideal, $S \neq \emptyset$. We partially order S by inclusion.

Let $\{I_\alpha\}_{\alpha \in A}$ be a totally ordered set of proper ideals in R. To write down an upper bound for these ideals in S, it is natural to try their union $I = \bigcup_{\alpha \in A} I_\alpha$. As with subgroups in the proof of Theorem 2.3, a union of ideals is *not* usually an ideal (try $2\mathbf{Z} \cup 3\mathbf{Z}$), but since we are dealing with a union of a totally ordered set of ideals, the union turns out to be an ideal. *Make sure you can explain why.*

If x and y are in I then $x \in I_\alpha$ and $y \in I_\beta$ for two of the ideals I_α and I_β. Since this set of ideals is totally ordered, $I_\alpha \subset I_\beta$ or $I_\beta \subset I_\alpha$. Without loss of generality, $I_\alpha \subset I_\beta$. Therefore x and y are in I_β, so $x \pm y \in I_\beta \subset I$. Hence I is an additive subgroup of R. The reader can check $rx \in I$ for $r \in R$ and $x \in I$, so I is an ideal in R.

Because I contains every I_α, I is an upper bound on the totally ordered subset $\{I_\alpha\}_{\alpha \in A}$ *provided* it is actually in S: is I a proper ideal? Well, if I is not a proper ideal then $1 \in I$. Since I is the union of the I_α's, we must have $1 \in I_\alpha$ for some α, but then I_α is not a proper ideal. That is a contradiction, so $1 \notin I$. Thus $I \in S$ and we have shown every totally ordered subset of S has an upper bound in S.

By Zorn's lemma S contains a maximal element. This maximal element is a proper ideal of R that is maximal with respect to inclusion among all proper ideals (not properly contained in another proper ideal of R). That means it is a maximal ideal of R. □

Corollary 3.2. *Every proper ideal in a nonzero commutative ring is contained in a maximal ideal.*

Proof. Let R be the ring and I be a proper ideal in R. The quotient ring R/I is nonzero, so it contains a maximal ideal by Theorem 3.1. The inverse image of this ideal under the natural reduction map $R \to R/I$ is a maximal ideal of R that contains I. □

It is crucial in the proof of Theorem 3.1 to have the multiplicative identity 1 available, which lies in no proper ideal. For instance, the analogue of Theorem 3.1 for groups can fail: the additive group $\mathbf{Q}$ contains no maximal proper subgroups, as we saw at the start of Section 2. The importance of Theorem 3.1 in the foundations of commutative algebra is one reason that rings should always have a multiplicative identity, at least if you are interested in areas of math that depend on commutative algebra, *e.g.*, number theory and algebraic geometry.

The following important theorem concerns nilpotent elements. In a commutative ring, an element r is called *nilpotent* if $r^n = 0$ for some $n \geq 1$.

Theorem 3.3 (Krull)**.** *The intersection of all prime ideals in a nonzero commutative ring is the set of nilpotent elements in the ring.*

This is striking: if we know an element in a nonzero commutative ring lies in every prime ideal, then a power of it is 0. In the ring $\mathbf{Z}$ the result is obvious, since the prime ideals are (0) and (p) for prime numbers p, and the intersection is obviously $\{0\}$, which is the only nilpotent integer. A somewhat more interesting example is the ring $\mathbf{Z}/(12)$. Its prime ideals are $(2)/(12)$ and $(3)/(12)$ (not $(0)/(12)$!), whose intersection is $(6)/(12) = \{0, 6 \bmod 12\}$, which is also (by inspection) all nilpotent elements of $\mathbf{Z}/(12)$.

Proof. Let R be a nonzero commutative ring. If $r \in R$ is nilpotent then $r^n = 0$ for some $n \geq 1$. For each prime ideal P of R, $r^n \in P$, so $r \in P$ since P is a prime ideal. Thus every nilpotent element of R is in the intersection of all prime ideals of R.

Now we want to show the intersection of all prime ideals of R consists only of nilpotent elements: if $r \in P$ for all prime ideals P then $r^n = 0$ for some $n \geq 1$. How could this be shown? We will *not* try to prove $r^n = 0$ for some n directly, but rather prove the contrapositive statement: if $r \in R$ is not nilpotent then some prime ideal of R does not contain r. So an element of R lying in all prime ideals of R must be nilpotent.

Since r is not nilpotent, $r^n \neq 0$ for every $n \geq 1$. Consider the set S of all ideals I in R that don't contain a positive power of r:

$$I \in S \Longleftrightarrow \{r^n : n \geq 1\} \cap I = \emptyset.$$

The zero ideal (0) doesn't meet $\{r^n : n \geq 1\}$, because r is not nilpotent, so S is nonempty. We partially order S by inclusion. After checking the conditions for Zorn's lemma can be applied to S, we will show that a maximal element of S is a prime ideal.[3] Since none of the ideals in S contain r, we will have found a prime ideal of R not containing r.

Let $\{I_\alpha\}_{\alpha \in A}$ be a totally ordered set of ideals in S. Its union I is an ideal (same proof as that in Theorem 3.1 – be sure you can write it up!). Since no I_α contains a positive power of r, their union I does not contain a positive power of r either. Thus $I \in S$, and since $I_\alpha \subset I$ for all α, I is an upper bound in S for the set of ideals I_α's. Thus every totally ordered subset of S contains an upper bound in S.

By Zorn's lemma, S has a maximal element. Call it P, so P is an ideal in R that does not contain a positive power of r and is maximal for this property (with respect to inclusion). We write a maximal element of S as P because we're going to show P is a prime ideal. The ideal P is proper since $r \notin P$. Suppose x and y are in R and $xy \in P$. To prove $x \in P$ or $y \in P$, assume otherwise. Then the ideals $(x) + P$ and $(y) + P$ are both strictly larger than P, so they can't lie in S. That means we have $r^m \in (x) + P$ and $r^n \in (y) + P$ for some positive integers m and n:

$$r^m = ax + p_1, \quad r^n = by + p_2$$

where p_1 and p_2 are in P and a and b are in R. Now multiply:

$$r^{m+n} = abxy + axp_2 + byp_1 + p_1p_2.$$

[3] Notice we are *not* defining S to be a set of prime ideals, but only a set of ideals with a disjointness property. That a maximal element of S is a prime ideal in R is going to be proved using maximality.

Since P is an ideal and $p_1, p_2, xy \in P$, the right side is in P. Then $r^{m+n} \in P$, which contradicts P being disjoint from $\{r^k : k \geq 1\}$ (because $P \in S$)[4]. Hence $x \in P$ or $y \in P$, so P is prime. By construction P contains no positive power of r, so in particular $r \notin P$. □

Remark 3.4. Although the ideal P in the proof above is maximal with respect to inclusion among the ideals in R that are disjoint from $\{r, r^2, r^3, \ldots\}$, P *need not* be a maximal ideal of R: a maximal ideal in a ring is maximal with respect to inclusion among *all* proper ideals of the ring, while the proof above used Zorn's lemma on a set of ideals in R that is usually not all proper ideals. (For example, if R is the set of fractions with odd denominator and $r = 2$ then $S = \{(0)\}$, so $P = (0)$, which is not a maximal ideal of R.) Quite generally, if S is a partially ordered set and S' is a subset of S, a maximal element of S' need not be a maximal element of S. Make sure you understand that!

The intersection of all prime ideals is called the *nilradical* (because it is the nilpotent elements) and the intersection of all maximal ideals is called the *Jacobson radical* (because Nathan Jacobson studied it). Since every prime ideal is in a maximal ideal, the intersection of the maximal ideals contains the intersection of the prime ideals and the containment can be strict, although examples of such rings are not typically seen in a first abstract algebra course. If every prime ideal is the intersection of the maximal ideals containing it then the ring is called a *Jacobson ring* and Theorem 3.3 with "prime ideal" replaced by "maximal ideal" is true for Jacobson rings. An example of a non-Jacobson ring is an integral domain R that is not a field and has one maximal ideal M: the nilradical of R is $\{0\}$ since R is an integral domain and the Jacobson radical of R is M, which is not $\{0\}$ since R is not a field. Such rings show up all over the place in commutative algebra and algebraic geometry, and include the ring of formal power series $F[[x]]$ where F is a field and the ring of fractions a/b where $a, b \in \mathbf{Z}$ and b is odd ("localization of $\mathbf{Z}$ at the prime 2").

Exercise. Use Zorn's lemma to prove that if D is a nonempty multiplicatively closed subset of a nonzero commutative ring R such that $0 \notin D$, then there is an ideal in R that is maximal with respect to being disjoint from D and such an ideal is a prime ideal. This result is due to Krull [19, Lemma, p. 732], and it implies Theorem 3.1 when $D = \{1\}$ and it implies Theorem 3.3 when $D = \{1, r, r^2, r^3, \ldots\}$ for a non-nilpotent $r \in R$.

Corollary 3.5. *For each proper ideal I in a nonzero commutative ring R,*

$$\bigcap_{P \supset I} P = \{x \in R : x^n \in I \text{ for some } n \geq 1\},$$

where the intersection runs over the prime ideals of R that contain I.

Proof. If $x^n \in I$ for some $n \geq 1$, then for each prime ideal $P \supset I$ we have $x^n \in P$, so $x \in P$. Conversely, suppose x is in every prime ideal of R containing I. (There are prime ideals containing I since there is a maximal ideal of R containing I.) The natural map $R \to R/I$ identifies the prime ideals in R that contain I with the prime ideals of R/I, so $\overline{x}$ is in every prime ideal of R/I. Therefore by Theorem 3.3, $\overline{x}$ is nilpotent in R/I. This means $\overline{x}^n = \overline{0}$ for some $n \geq 1$, so $x^n \in I$. □

[4] If we had defined S to be the ideals in R that don't include r, rather than no positive power of r, then at this point we'd be stuck: we'd have $r = ax + p_1$ and $r = bx + p_2$, so then multiplying gives $r^2 \in P$, but that would not be a contradiction since we didn't require ideals of S to avoid containing r^2. Seeing this difficulty is a motivation for the definition of S in the proof.

We used Zorn's lemma to prove Theorem 3.3 since the point of this handout is to see enough applications of Zorn's lemma that the basic principle behind its use becomes transparent, but it turns out that Theorem 3.3 can also be proved using Corollary 3.2 about proper ideals lying in a maximal ideal. The trick is to use a maximal ideal not in R itself but in the polynomial ring $R[X]$. Here is that alternate proof of Theorem 3.3.

Proof. We will only address one direction: a non-nilpotent element r in a ring R lies outside some prime ideal. (The other direction is easy; see the first paragraph in the first proof of Theorem 3.3.) We will create such a prime ideal as the kernel of a homomorphism out of R and r won't be in the kernel.

In $R[X]$, $rX-1$ is not a unit. Indeed, if we did have $(rX-1)(c_nX^n+\cdots+c_1X+c_0)=1$ then equating coefficients of like powers of X on both sides shows

$$-c_0=1, \quad c_0r-c_1=0, \quad c_1r-c_2=0, \ \ldots, \quad c_{n-1}r-c_n=0, \quad rc_n=0.$$

Thus $c_0=-1$, $c_i=c_{i-1}r$ for $1\leq i\leq n$, and $rc_n=0$. So $c_i=-r^i$ for $1\leq i\leq n$. Thus $0=rc_n=-r^{n+1}$, but r is not nilpotent. So $rX-1$ is not a unit in $R[X]$. (If r *were* nilpotent, say $r^n=0$, then $rX-1$ is a unit in $R[X]$ with inverse $-(1+rX+r^2X^2+\cdots+r^{n-1}X^{n-1})$.)

The ideal $(rX-1)$ in $R[X]$ is proper, since $rX-1$ is not a unit, so this ideal lies inside a maximal ideal M of $R[X]$ by Corollary 3.2. Now consider the composite homomorphism $R\to R[X]\to R[X]/M$, where the first map is inclusion and the second is reduction. Since the target is a field, the kernel in R is a prime ideal (the quotient of R by the kernel embeds into the field $R[X]/M$ and thus must be an integral domain, as subrings of fields are integral domains). Call the kernel P. Since $rX\equiv 1 \bmod M$, r is not in the kernel. Thus P is a prime ideal in R not containing r. We're done. □

This proof can be streamlined if you know about localization of commutative rings. (If you don't, skip this paragraph.) For a commutative ring R and multiplicative subset S of R, the localization R_S is 0 if and only if $0\in S$. Taking $S=\{1,r,r^2,\ldots\}$ where r is not nilpotent, R_S is usually written as $R[1/r]$. (It is isomorphic to the ring $R[X]/(rX-1)$ used in the above proof.) Since $R[1/r]\neq\{0\}$, $R[1/r]$ contains a maximal ideal M and the image of r in $R[1/r]$ is a unit and thus is not in M. The kernel of the composite homomorphism $R\to R[1/r]\to R[1/r]/M$ is a prime ideal for the same reason given in the proof above.

We now leave maximal ideals and nilpotent elements, turning our attention to an interesting theorem of Cohen [8, Theorem 2] about finitely generated ideals.

Theorem 3.6 (Cohen). *If every prime ideal in a commutative ring is finitely generated then every ideal in the ring is finitely generated.*

Proof. We will prove the contrapositive: if there is some ideal in the ring that is not finitely generated then there is a prime ideal in the ring that is not finitely generated. We will find this prime ideal as an ideal maximal with respect to inclusion for the property of not being finitely generated. (Here again we should stress, as in Remark 3.4, that such prime ideals need not be actual maximal ideals in the ring. They are only created as being maximal among non-finitely generated ideals.)

Let S be the collection of all non-finitely generated ideals, so $S\neq\emptyset$ by *assumption*. We partially order S by inclusion. For a totally ordered subset of ideals $\{I_\alpha\}_{\alpha\in A}$ in S, its union I is an ideal containing each I_α. (That I is an ideal follows by the same argument as in the proof of Theorem 3.1.) To know I is an upper bound on the I_α's in S we have to show I is not finitely generated. Well, if I were finitely generated, say $I=(r_1,\ldots,r_k)$, then each of the generators is in some I_α and by total ordering on the ideals, these hypothetical finitely many generators of I are all in a common I_α (Lemma 1.9). But then I lies inside that I_α, so I equals that I_α, which shows that I_α is finitely generated. This is a contradiction, so I is not finitely generated.

Now we can apply Zorn's lemma: there exists a non-finitely generated ideal that is maximal with respect to inclusion among the non-finitely generated ideals. Call such an ideal P. We are going to prove P is a prime ideal. It is certainly a proper ideal. Suppose $xy \in P$ with $x \notin P$ and $y \notin P$. Then $(x) + P$ is an ideal properly containing P, so $(x) + P \notin S$. Therefore this ideal is finitely generated:

$$(x) + P = (r_1, \dots, r_k).$$

Write $r_i = c_i x + p_i$ for $i = 1, 2, \dots, k$, where $c_i \in R$ and $p_i \in P$. (We have no right to expect the p_i's generate P. They only occur in the expressions for the r_i's.) Then every r_i is in the ideal $(x, p_1, \dots, p_k)$, and conversely each p_i is in the ideal $P \subset (x) + P = (r_1, \dots, r_k)$, so

$$(x) + P = (x, p_1, \dots, p_k).$$

If we are given $p \in P$, then since $P \subset (x) + P$ we can write

$$p = cx + a_1 p_1 + \cdots + a_k p_k \tag{3.1}$$

with c and the a_i's all in R. Then $cx = p - \sum a_i p_i \in P$, so c lies in the ideal $J = \{r \in R : rx \in P\}$. Obviously $P \subset J$. Since $xy \in P$ and $y \notin P$, J contains y and therefore J strictly contains P. Thus by maximality of P among all the non-finitely generated ideals in R, J is finitely generated. By (3.1), $p \in xJ + \sum_{i=1}^k Rp_i$, so $P \subset xJ + \sum_{i=1}^k Rp_i$. The reverse inclusion is easy (by the definition of J), so

$$P = xJ + \sum_{i=1}^{k} Rp_i.$$

Since J is finitely generated, this shows P is finitely generated, a contradiction. Hence $x \in P$ or $y \in P$, so P is a prime ideal. □

Exercise. Use Zorn's lemma to prove an analogue of Theorem 3.6 for principal ideals: if every prime ideal in a commutative ring is principal than all ideals are principal. (It is false that if every prime ideal has at most 2 generators then all ideals have at most 2 generators, *e.g.*, in $\mathbf{C}[X, Y]$ the prime ideals have 1 or 2 generators but the ideal (X^2, XY, Y^2) can't be generated by 2 elements. It is also false that if every maximal ideal is principal then all ideals are principal: see `https://mathoverflow.net/questions/81011`.)

We now generalize Theorem 3.6 to modules in place of rings. When M is an R-module and $\mathfrak{a}$ is an ideal of R, let $\mathfrak{a}M$ denote the submodule of M that is spanned by all finite products am for $a \in \mathfrak{a}$ and $m \in M$. That is, $\mathfrak{a}M$ is the set of all finite sums $\sum_{i=1}^n a_i m_i$ with $n \geq 1$, $a_i \in \mathfrak{a}$ and $m_i \in M$. Check this is a submodule of M. (We need to use finite sums since the set of products am with $a \in \mathfrak{a}$ and $m \in M$ is usually not a submodule since it's not additively closed.)

Theorem 3.7 (Jothilingam [16]). *Let M be a finitely generated R-module. If every submodule of the form $\mathfrak{p}M$ with prime $\mathfrak{p}$ is finitely generated then every submodule of M is finitely generated.*

When $M = R$, this is Theorem 3.6. However, the proof of Theorem 3.7 has some additional aspects at the end that don't occur in Theorem 3.6, so we are proving the two theorems separately. Our proof is based on an argument of Naghipour [23].

Proof. We will prove the contrapositive: if M has a submodule that is not finitely generated then it has a submodule of the form $\mathfrak{p}M$ with prime $\mathfrak{p}$ that is not finitely generated.

Let S be the set of submodules of M that are not finitely generated, so $S \neq \emptyset$ by *assumption*. Note $M \notin S$. Partially order S by inclusion. By the same kind of argument as in the proof of Theorem 3.6, every totally ordered subset of S has an upper bound in S. (Check the details!) Therefore we can apply Zorn's lemma: S contains a maximal element. Call one of them N. That is, N is a submodule of M that is not finitely generated and (this is the key point) each submodule of M that properly contains N is finitely generated. Note $N \neq M$.

Will we show $N = \mathfrak{p}M$ for some prime ideal $\mathfrak{p}$ of R? No! There is so little control over the maximal elements coming from Zorn's lemma that we can't expect this. Instead we will show that

(1) $\mathfrak{p} := \mathrm{Ann}_R(M/N) = \{r \in R : rM \subset N\}$ is a prime ideal of R,
(2) $\mathfrak{p}M$ is not finitely generated.

To show $\mathfrak{p}$ is prime, first we note $\mathfrak{p} \neq R$ since $N \subsetneqq M$. If $\mathfrak{p}$ is not prime then there are x and y in R with $xy \in \mathfrak{p}$ but x and y are not in $\mathfrak{p}$. From the definition of $\mathfrak{p}$, these conditions on x and y mean

$$xyM \subset N, \quad xM \not\subset N, \quad yM \not\subset N.$$

Thus $xM + N$ properly contains N, so $xM + N$ is finitely generated. Let a finite spanning set of $xM + N$ be $xm_i + n_i$ $(i = 1, \ldots, k)$. (warning: $n_1, \ldots, n_k$ do not span N, as N is not finitely generated.) Then

$$xM + N = \sum_{i=1}^{k} Rxm_i + \sum_{i=1}^{k} Rn_i.$$

(Just check a spanning set of the module on each side is in the other side.) For all $n \in N \subset xM + N$,

$$\begin{aligned} n &= r_1xm_1 + \cdots + r_kxm_k + r_1'n_1 + \cdots + r_k'n_k \\ &= x(r_1m_1 + \cdots + r_km_k) + r_1'n_1 + \cdots + r_k'n_k \end{aligned} \tag{3.2}$$

where $r_i, r_i' \in R$. Thus $x(r_1m_1 + \cdots + r_km_k) \in N$, so $r_1m_1 + \cdots + r_km_k$ lies in

$$L := \{m \in M : xM \subset N\},$$

which is a submodule of M. Note

$$N \subset L, \quad yM \subset L, \quad yM \not\subset N.$$

Therefore N is a proper subset of L, so L is finitely generated. By (3.2),

$$n \in xL + \sum_{i=1}^{k} Rn_i,$$

so

$$N \subset xL + \sum_{i=1}^{k} Rn_i.$$

The reverse inclusion is straightforward (use the definition of L), so

$$N = xL + \sum_{i=1}^{k} Rn_i.$$

The right side is finitely generated, which is a contradiction since N is not finitely generated. Thus $\mathfrak{p}$ is a prime ideal in R.

(At this point, if we had been taking $M = R$ as in Theorem 3.6, then N would be an ideal and $\operatorname{Ann}_R(M/N) = \operatorname{Ann}_R(R/N)$ would equal N, so $N = \mathfrak{p}$ would be prime and we would have finished proving Theorem 3.6.)

It remains to show $\mathfrak{p}M$ is not finitely generated. What we will do is show $N = \mathfrak{p}M + Q$ for some finitely generated R-module Q. Then, since N is not finitely generated, $\mathfrak{p}M$ can't be finitely generated either.

Since, by hypothesis, M is finitely generated, write $M = Re_1 + \cdots + Re_\ell$. (We are not assuming the e_i's are linearly independent, just that they form a spanning set.) Then M/N is spanned over R by the reductions $\overline{e}_1, \ldots, \overline{e}_k$, so

$$\mathfrak{p} = \operatorname{Ann}_R(M/N) = \bigcap_{i=1}^{\ell} \operatorname{Ann}_R(R\overline{e}_i).$$

The ideal $\mathfrak{p}$ is inside each $\operatorname{Ann}_R(R\overline{e}_i)$, but in fact it must equal one of these: if not then each $\operatorname{Ann}_R(R\overline{e}_i)$ contains an element r_i outside $\mathfrak{p}$, but then the product of those r_i's (over all i) is an element outside of $\mathfrak{p}$ (because $\mathfrak{p}$ is prime) while at the same time the product of the r_i's kills each $\overline{e}_i$, so this product is in each annihilator and hence is in $\mathfrak{p}$. This is absurd, so $\mathfrak{p}$ is the annihilator of some $R\overline{e}_i$. Without loss of generality, $\mathfrak{p} = \operatorname{Ann}_R(R\overline{e}_1) = \{r \in R : re_1 \subset N\}$.

Since $\mathfrak{p} \neq R$, also $\overline{e}_1 \neq 0$ in M/N, and thus $e_1 \notin N$, so $Re_1 \not\subset N$. Therefore $Re_1 + N$ properly contains N so it is finitely generated, say by $r_je_1 + n_j$ ($j = 1, \ldots, d$). For $n \in N \subset Re_1 + N$, write

$$n = \sum_{j=1}^{d} a_j(r_je_1 + n_j) = \left(\sum_{j=1}^{d} a_jr_j\right) e_1 + \sum_{j=1}^{d} a_jn_j,$$

with $a_j \in R$. The coefficient of e_1 scales e_1 into N by this equation, so the coefficient of e_1 is in $\mathfrak{p}$. Thus

$$N \subset \mathfrak{p}e_1 + \sum_{j=1}^{d} Rn_j.$$

The reverse inclusion is easy, so

$$N = \mathfrak{p}e_1 + \sum_{j=1}^{d} Rn_j \subset \mathfrak{p}M + \sum_{j=1}^{d} Rn_j \subset N + N = N.$$

Thus $N = \mathfrak{p}M + \sum_{j=1}^{d} Rn_j$, so N not being finitely generated forces $\mathfrak{p}M$ not to be finitely generated, which is what we wanted to show. □

Remark 3.8. In Theorems 3.3, 3.6, and 3.7, an ideal built with Zorn's lemma turns out to be a prime ideal. Theorem 3.1 is also such a result, since maximal ideals in a commutative ring are prime ideals. Lam and Reyes [20] describe general conditions under which this kind of phenomenon occurs.

4. Applications to bases of vector spaces

We want to use Zorn's lemma to prove an arbitrary nonzero vector space has a basis. Let's first make sure we know what the label "basis" means when we are dealing with vector spaces that may turn out to be infinite-dimensional. For a nonzero vector space V over a field F, a *basis* of V is a subset $\mathcal{B}$ of V that is linearly independent (*i.e.*, no finite subset of $\mathcal{B}$ has a nontrivial F-linear relation) and spans V (*i.e.*, every element of V is an F-linear combination of finitely many elements of $\mathcal{B}$).

Even if a basis is infinite, finiteness assumptions are built into linear independence and spanning sets: linear independence involves a finite linear combination equal to 0, and spanning sets involve finitely many vectors at a time. In analysis, infinite linear combinations occur when a topology has been introduced on the vector space. Such a topological basis is not covered by the use of the label "basis" here. Our more algebraically-oriented concept of a basis, always using finite linear combinations, is called a Hamel basis.

Theorem 4.1 (Hausdorff)**.** *Every nonzero vector space contains a basis.*

This was proved by Hamel [12] in the special case of $\mathbf{R}$ as a vector space over $\mathbf{Q}$ in 1905. The result in full generality is due to Hausdorff [13, p. 295] in 1932.

Proof. The idea is that a basis can be constructed as a maximal linearly independent set, and this maximal set will be found with Zorn's lemma.

Let V be a nonzero vector space and let S be the set of linearly independent sets in V. For instance, a single nonzero $v \in V$ is a linearly independent set, so $\{v\} \in S$. Thus $S \neq \emptyset$.

For two linearly independent sets L and L' in V, declare $L \leq L'$ if $L \subset L'$. This is the partial ordering on S by inclusion. It is easy to see that every subset of a linearly independent set is also a linearly independent set, so if $L \in S$ then every subset of L is also in S.

Assume $\{L_\alpha\}_{\alpha \in A}$ is a totally ordered subset of S. That is, every L_α is a linearly independent set in V and for each L_α and L_β in our subset we have $L_\alpha \subset L_\beta$ or $L_\beta \subset L_\alpha$. An upper bound for the L_α's in S is the union $L = \bigcup_{\alpha \in A} L_\alpha$. Well, we need to check L is really a linearly independent set (so $L \in S$); once that is settled then L is an upper bound in S since $L_\alpha \subset L$ for all $\alpha \in A$.

Pick a finite set of vectors $v_1, \ldots, v_n$ in L. We must show they are linearly independent. Each v_k is in some L_α, say $v_1 \in L_{\alpha_1}, \ldots, v_n \in L_{\alpha_n}$. Since the L_α's are totally ordered, one of the sets $L_{\alpha_1}, \ldots, L_{\alpha_n}$ contains the others (Lemma 1.9). That means $v_1, \ldots, v_n$ are all in a common L_α, so they are linearly independent.

Zorn's lemma now tells us that S contains a maximal element: there is a linearly independent set $\mathcal{B}$ in V that is not contained in a larger linearly independent set in V. We will show $\mathcal{B}$ spans V, so it is a basis.

Let W be the span of $\mathcal{B}$. That means W is the set of all finite F-linear combinations $\sum_{i=1}^k c_i v_i$ with $k \geq 1$, $c_i \in F$, and $v_i \in \mathcal{B}$. If $\mathcal{B}$ does not span V then $W \neq V$, so we can pick $v \in V$ with $v \notin W$. Then $\mathcal{B}$ is a proper subset of $\mathcal{B} \cup \{v\}$. We will show $\mathcal{B} \cup \{v\}$ is linearly independent, which contradicts the maximality of $\mathcal{B}$ and thus proves $W = V$.

To prove $\mathcal{B} \cup \{v\}$ is linearly independent, assume otherwise: there is an expression

$$\sum_{i=1}^{k} c_i v_i = 0 \tag{4.1}$$

where the coefficients are not all 0 and the v_i's are taken from $\mathcal{B} \cup \{v\}$. Since the elements of $\mathcal{B}$ are linearly independent, one of the v_i's with a nonzero coefficient must be v. We can re-index and suppose $v_k = v$, so $c_k \neq 0$. We must have $k \geq 2$, since otherwise $c_1 v = 0$, which is impossible since $v \neq 0$ and the coefficient of v is nonzero. Then

$$c_k v = -\sum_{i=1}^{k-1} c_i v_i.$$

Multiplying both sides by $1/c_k$,

$$v = \sum_{i=1}^{k-1} \left(-\frac{c_i}{c_k}\right) v_i,$$

which shows $v \in W$. But $v \notin W$, so $\mathcal{B} \cup \{v\}$ is a linearly independent set. □

Corollary 4.2. *Every linearly independent subset of a nonzero vector space V can be extended to a basis of V. In particular, every subspace W of V is a direct summand: $V = W \oplus U$ for some subspace U of V.*

Proof. Let $\mathcal{L}$ be a linearly independent subset of V. A basis of V containing $\mathcal{L}$ will be found as a maximal linearly independent subset containing $\mathcal{L}$.

Take S to be the set of linearly independent sets in V that contain $\mathcal{L}$. For instance, $\mathcal{L} \in S$, so $S \neq \emptyset$. The same argument as in the proof of Theorem 4.1 shows every totally ordered subset of S has an upper bound. (If the L_α's are linearly independent sets in V that each contain $\mathcal{L}$ then their union L also contains $\mathcal{L}$, and $L \in S$ because the L_α's are totally ordered, by a kind of argument we've made before.)

By Zorn's lemma there is a maximal element of S. This is a linearly independent set in V that contains $\mathcal{L}$ and is maximal with respect to inclusion among all linearly independent sets in S containing $\mathcal{L}$. The proof that a maximal element of S is a basis of V follows just as in the proof of Theorem 4.1.

To prove every subspace $W \subset V$ is a direct summand, let $\mathcal{L}$ be a basis of W. There is a basis $\mathcal{B}$ of V containing $\mathcal{L}$. Let U be the span of the complement $\mathcal{B} - \mathcal{L}$. It is left to the reader to show $V = W + U$ and $W \cap U = \{0\}$, so $V = W \oplus U$. □

Corollary 4.3. *Every spanning set of a nonzero vector space V contains a basis of V.*

Proof. Let $\mathcal{S}$ be a spanning set of V. Consider the set of linearly independent subsets of $\mathcal{S}$. This is a nonempty set, as $\{v\}$ is linearly independent for each $v \in \mathcal{S}$. Partially order the set of linearly independent subsets of $\mathcal{S}$ by inclusion. If $\{L_i\}$ is a totally ordered subset then $\bigcup_i L_i$ is a linearly independent subset of $\mathcal{S}$ and an upper bound on the L_i's. So by Zorn's lemma there is a maximal element $\mathcal{B}$: a linearly independent subset of $\mathcal{S}$ that is maximal with respect to inclusion. We will show $\mathcal{B}$ is a spanning set for V so it is a basis. Because $\mathcal{S}$ spans V, it is enough to show every element of $\mathcal{S}$ is in the span of $\mathcal{B}$ to know V is spanned by $\mathcal{B}$. If some $v \in \mathcal{S}$ is not in the span of $\mathcal{B}$ then $\mathcal{B} \cup \{v\}$ is a linearly independent set and it is a subset of $\mathcal{S}$ that strictly contains $\mathcal{B}$. This contradicts the maximality of $\mathcal{B}$ in $\mathcal{S}$. □

Remark 4.4. To find a basis of V inside a spanning set $\mathcal{S}$, a natural first idea might be to find a minimal spanning set of V inside of $\mathcal{S}$ rather than a maximal linearly independent subset. The minimality of a spanning set would force its linear independence and thus give us a basis. It is obvious how to use Zorn's lemma here: consider the set of all spanning sets of V inside $\mathcal{S}$, and partially order it by reverse inclusion. If $\{S_i\}$ is a totally ordered subset then the intersection $\bigcap_i S_i$ should be an upper bound on all the S_i's (we're using *reverse* inclusion, so an upper bound is a spanning set contained *in* every S_i), and then Zorn's lemma gives us maximal elements, which will be minimal spanning sets. But there's a problem: how do you prove $\bigcap_i S_i$ is a spanning set? You can't; it's not generally true. For example, let $V = \mathbf{Q}$ as a $\mathbf{Q}$-vector space, enumerate the rationals as $\{r_1, r_2, r_3, \dots\}$ and let S_i equal $\mathbf{Q}$ with the first i rationals removed. Each S_i is a spanning set of $\mathbf{Q}$ as a $\mathbf{Q}$-vector space, and the S_i's are totally ordered by reverse inclusion, but their intersection is *empty*. This is an instructive case where it seems clear how Zorn's lemma should work, but it doesn't work!

Here are some amusing corollaries of the existence of bases in an arbitrary (especially infinite-dimensional) vector space.

Corollary 4.5 (Hamel)**.** *There is a function $f\colon \mathbf{R} \to \mathbf{R}$ satisfying $f(x+y) = f(x) + f(y)$ for all $x, y \in \mathbf{R}$ that is not of the form $f(x) = cx$ for some $c \in \mathbf{R}$.*

Proof. Since $\mathbf{Q} \subset \mathbf{R}$, we can think of $\mathbf{R}$ as a $\mathbf{Q}$-vector space. In $\mathbf{R}$, the numbers 1 and $\sqrt{2}$ are linearly independent over $\mathbf{Q}$ since $\sqrt{2}$ is irrational. The $\mathbf{Q}$-linearly independent subset $\{1, \sqrt{2}\}$ of $\mathbf{R}$ can be extended to a $\mathbf{Q}$-basis $\{e_i\}$ of $\mathbf{R}$ by Corollary 4.2. Define a bijection f of the $\mathbf{Q}$-basis by $f(1) = \sqrt{2}$, $f(\sqrt{2}) = 1$, and $f(e_i) = e_i$ for $e_i \neq 1$ or $\sqrt{2}$. This bijection extends to a $\mathbf{Q}$-linear map $f\colon \mathbf{R} \to \mathbf{R}$ by $f(\sum r_i e_i) = \sum r_i f(e_i)$ on finite linear combinations of the $\mathbf{Q}$-basis, where $r_i \in \mathbf{Q}$. Since f is $\mathbf{Q}$-linear, it is an additive function $\mathbf{R} \to \mathbf{R}$. Since $f(1) = \sqrt{2}$ and $f(\sqrt{2}) = 1$, there is no c such that $f(x) = cx$ for all $x \in \mathbf{R}$ (for $x = 1$ we'd need $c = \sqrt{2}$ and for $x = \sqrt{2}$ we'd need $c = 1/\sqrt{2}$). □

In this proof, 1 and $\sqrt{2}$ can be replaced by two elements of an arbitrary $\mathbf{Q}$-basis of $\mathbf{R}$. There is no known way to describe a function f as in Corollary 4.5 that avoids using a $\mathbf{Q}$-basis of $\mathbf{R}$ and there is no concrete formula for a $\mathbf{Q}$-basis of $\mathbf{R}$: it exists purely by Zorn's lemma.

Corollary 4.6. *As abelian groups, $\mathbf{R}^n$ is isomorphic to $\mathbf{R}$ for each $n \geq 1$.*

Proof. Let $\{x_i\}_{i \in I}$ be a $\mathbf{Q}$-basis of $\mathbf{R}$. Since $\mathbf{R}^n$ is an $\mathbf{R}$-vector space with the obvious basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$, $\mathbf{R}^n$ as a $\mathbf{Q}$-vector space has basis $\{x_i \mathbf{e}_j\}$ where $i \in I$ and $j \in \{1, \dots, n\}$. Since I is infinite, $I \times \{1, \dots, n\}$ has the same cardinality as I. Therefore the $\mathbf{Q}$-bases of $\mathbf{R}^n$ and $\mathbf{R}$ have the same cardinality.

Since $I \times \{1, \dots, n\}$ and I are in bijection, we can write a $\mathbf{Q}$-basis of $\mathbf{R}^n$ using index set I, say $\{\mathbf{x}_i\}$. Then a group isomorphism $\mathbf{R} \to \mathbf{R}^n$ is obtained by sending x_i to $\mathbf{x}_i$ and extending this by $\mathbf{Q}$-linearity: $\sum_i r_i x_i \mapsto \sum_i r_i \mathbf{x}_i$ for all finite $\mathbf{Q}$-linear combinations of the $\mathbf{Q}$-basis of $\mathbf{R}$. □

Corollary 4.7. *If G is a group with more than two elements then G has a nontrivial automorphism.*

Proof. If G is nonabelian then some element of G is not in the center, so conjugation by that element is a nontrivial automorphism of G. If G is abelian then inversion (sending each element to its inverse) is an automorphism, and it is nontrivial unless every element is its own inverse. If G is abelian and every element is its own inverse then every element is killed by 2 ($x = -x \Rightarrow 2x = 0$), so G is a vector space over $\mathbf{Z}/(2)$. Let $\{e_i\}_{i \in I}$ be a basis of G over $\mathbf{Z}/(2)$. If there is more than one basis element, then exchanging two basis elements while fixing the rest extends to an automorphism of G. The only case remaining is a $\mathbf{Z}/(2)$-vector space with a basis of size at most 1. Such groups are trivial or cyclic of order 2, and their only automorphism is the identity. □

In the proofs of Theorem 4.1 and Corollaries 4.2 and 4.3, we did not use commutativity of the coefficient field under multiplication except at the end of the proof of Theorem 4.1, where we wrote c_i/c_k. If we write that more carefully as $c_k^{-1}c_i$ then the proof goes through when the coefficient ring is a division ring: multiply through on the left by c_k^{-1} to solve for v as a linear combination of $v_1, \ldots, v_{k-1}$. The proofs of Corollaries 4.2 and 4.3 also go through with a division ring as the coefficient ring: every nonzero vector space V over a division ring has a basis, every linearly independent subset of V can be extended to a basis, and every spanning set of V contains a basis.

Modules over commutative rings that are *not* fields need not have a basis. For example, a nonprincipal ideal in a commutative ring R is an R-module without a basis: if $a, b \in R$ are nonzero then they are R-linearly dependent since $b \cdot a + (-a) \cdot b = 0$, so a basis of an ideal as an R-module can't contain more than one element.[5] In a module that has a basis, a linearly independent subset need not extend to a basis and a submodule need not be a direct summand. For instance, the one-element linearly independent set $\{(4,6)\}$ in $\mathbf{Z}^2$ can't be extended to a basis of $\mathbf{Z}^2$ and the submodule $(2\mathbf{Z})^2$ is not a direct summand of $\mathbf{Z}^2$. The reason our proofs for vector spaces over fields don't carry over to $\mathbf{Z}$-modules is that nonzero integers generally don't have inverses in $\mathbf{Z}$.

5. Application to connected components

A subset C of a topological space X is called *connected* if, whenever $C \subset U \cup V$ for disjoint open subsets U and V of X, either $C \subset U$ or $C \subset V$. This means it is impossible to decompose C into two parts that lie in disjoint open subsets of X. As a trivial example, one-element subsets of X are connected. For a nontrivial example, it is fundamental property of $\mathbf{R}$ that the open interval $(0,1)$ is connected, and more generally all intervals (a,b), $(a,b]$, $[a,b)$, and $[a,b]$ with $a < b$ are connected. For $n \geq 1$, every open ball in $\mathbf{R}^n$ together with a subset of its boundary is connected.[6]

Lemma 5.1. *If $\{C_i\}_{i \in I}$ is a collection of nonempty connected subsets of a topological space X and whenever $i \neq j$ in I we have $C_i \cap C_j \neq \emptyset$, then $\bigcup_{i \in I} C_i$ is a connected subset of X.*

Proof. Set $C = \bigcup_{i \in I} C_i$ and suppose $C \subset U \cup V$ where U and V are disjoint open subsets of X. We want to prove $C \subset U$ or $C \subset V$.

Fixing $j \in I$, we have $C_j \subset U \cup V$, so $C_j \subset U$ or $C_j \subset V$ by connectedness of C_j. Without loss of generality, $C_j \subset U$.

For each $i \neq j$ we have $C_i \cap C_j \neq \emptyset$. Since $C_i \cap C_j \subset C_j \subset U$, C_i has nonempty intersection with U, so $C_i \subset U$ by connectedness of C_i. Thus $C_j \cup \bigcup_{i \in I - \{j\}} C_i \subset U$, so $C \subset U$. □

The set of nonzero real numbers $\mathbf{R}^\times = (-\infty, 0) \cup (0, \infty)$ is not connected, but its two parts $(-\infty, 0)$ and $(0, \infty)$ are each connected and maximally so: they do not belong to a larger connected subset of $\mathbf{R}^\times$. This idea generalizes to all topological spaces.

Definition 5.2. A *connected component* of a topological space X is a connected subset C that is maximal with respect to containment: C is connected, and if $C \subset \widetilde{C} \subset X$ with $\widetilde{C}$ connected, then $\widetilde{C} = C$.

[5] A nonzero ideal in R has a basis as an R-module if and only if it is a principal ideal with a generator that is not a zero divisor.

[6] It suffices to prove open balls in $\mathbf{R}^n$ are connected, since it can be shown that if C is a connected subset of X then each set lying between C and its closure $\overline{C}$ is connected.

Example 5.3. For $n \geq 1$, the topological space $\mathrm{GL}_n(\mathbf{R})$ has two connected components: invertible matrices with positive determinant and invertible matrices with negative determinant.

Example 5.4. The topological space $(\mathbf{R}^\times)^n$ has 2^n connected components, where the coordinates in each n-tuple have a fixed sign (positive or negative).

Here are two basic properties of connected components of a topological space X.

(1) If X is nonempty then its connected components are nonempty since one-element subsets of X are connected.
(2) Different connected components of X are disjoint: this is obvious if $X = \emptyset$, and if $X \neq \emptyset$ and C_1 and C_2 are connected components of X with $C_1 \cap C_2 \neq \emptyset$ then we will prove $C_1 = C_2$. The union $C_1 \cup C_2$ is connected since $C_1 \cap C_2 \neq \emptyset$ (Lemma 5.1), so from $C_1 \subset C_1 \cup C_2$ and maximality of connected components for containment we get $C_1 = C_1 \cup C_2$. Therefore $C_2 \subset C_1$, and changing the roles of C_1 and C_2 gives us $C_1 \subset C_2$, so $C_1 = C_2$.[7]

The following theorem about the existence of connected components is analogous to the existence of maximal ideals in nonzero commutative rings (Theorem 3.1).

Theorem 5.5. *Every topological space has a connected component.*

Proof. The theorem is obvious for the empty space, so let X be a nonempty topological space. Let S be the collection of nonempty connected subsets of X. We have $S \neq \emptyset$ since one-element subsets of X are in S. Partially order S by inclusion.

Let $\{C_\alpha\}_{\alpha \in A}$ be a totally ordered subset of S: each C_α is connected and for all α and β in A we have either $C_\alpha \subset C_\beta$ or $C_\beta \subset C_\alpha$. To write down an upper bound for the C_α's in S, it is natural to try their union $C := \bigcup_{\alpha \in A} C_\alpha$. As a set, C certainly contains all the C_α's, but is C connected? A union of connected subsets is *not* usually connected, such as $(0,1) \cup (2,3)$ in $\mathbf{R}$. However, for all α and β in A we have $C_\alpha \subset C_\beta$ or $C_\beta \subset C_\alpha$, so $C_\alpha \cap C_\beta \neq \emptyset$. Therefore all pairs of connected subsets in $\{C_\alpha\}_{\alpha \in A}$ have nonempty intersection, which implies the union $C := \bigcup_{\alpha \in A} C_\alpha$ is connected by Lemma 5.1, so C is an upper bound in S on the totally ordered subset $\{C_\alpha\}_{\alpha \in A}$. We have shown every totally ordered subset of S has an upper bound in S.

By Zorn's lemma S contains a maximal element M. Since M is maximal with respect to containment among connected subsets of X, M is a connected component of X. □

Actually, what we really want is a result analogous to Corollary 3.2.

Theorem 5.6. *Every nonempty connected subset of a topological space lies in a unique connected component. In particular, every point in a nonempty topological space lies in a unique connected component.*

Proof. Let N be a nonempty connected subset of a topological space X and let S be the collection of (nonempty) connected subsets of X that contain N. We have $S \neq \emptyset$ since $N \in S$.

For a totally ordered subset $\{C_\alpha\}_{\alpha \in A}$ of S the union $C := \bigcup_{\alpha \in A} C_\alpha$ contains N and C is connected since $\bigcap_{\alpha \in A} C_\alpha$ contains N and therefore is nonempty (Lemma 5.1). Thus every totally ordered subset of S has an upper bound in S.

[7] Since the closure of a connected subset is connected, by maximality a connected component has to be its own closure, so connected components are also closed.

By Zorn's lemma S contains a maximal element M. A connected subset of X containing M also contains N, so it must be M by maximality of M in S. That proves M is a connected component of X. If M' is a connected component of X containing N then $N \subset M \cap M'$, so $M \cap M' \neq \emptyset$. Thus $M' = M$ since different connected components are disjoint. □

Remark 5.7. The proof of Theorem 5.6 did not use the total ordering in the collection $\{C_\alpha\}_{\alpha \in A}$, which suggests this theorem is not as closely related to Zorn's lemma as some other applications of it such as Theorem 3.1. Indeed, the decomposition of X into connected components can be proved without using Zorn's lemma. Set $x \sim x'$ in X if x and x' lie in a common connected subset of X. Check this is an equivalence relation on X. Its equivalence classes are the connected components of X and this shows each element of X belongs to a (unique) connected component of X.

6. Equivalences and controversies with Zorn's lemma

We used Zorn's lemma to prove several existence theorems. Some of those results can be reversed: Zorn's lemma is equivalent (within Zermelo-Fraenkel set theory) to

- all nonzero vector spaces over all fields have a basis [5, 18],
- the existence of a maximal ideal in all nonzero commutative rings [1, 15],
- Theorem 2.4 [4, Theorem 2.1].

See [14, 26] for surveys on equivalents to, and consequences of, Zorn's lemma.

Remark 6.1. We saw in Theorem 4.1 that Zorn's lemma implies every nonzero vector space over every field has a basis. It is not obvious that Zorn's lemma is also implied by every nonzero vector space over every field having a basis. How does that work? The proof deduces the Axiom of Choice from the existence of bases. That the Axiom of Choice is equivalent to Zorn's lemma is a separate argument (we mentioned that in the later part of Section 1). To deduce the Axiom of Choice from the existence of bases uses an abstract vector space over a complicated field of rational functions. A brief summary of the argument is given by Noah Schweber in an answer at `https://math.stackexchange.com/questions/1650069/`. Zorn's lemma is *not* known to be a consequence of the existence of bases of vector spaces over fields of prior interest, like **Q** or **R**.

In topology, Tychonoff's theorem in its general form (allowing compact non-Hausdorff spaces) is equivalent to Zorn's lemma [14, Theorem 4.68], [17], but Tychonoff's theorem for products of compact Hausdorff spaces is equivalent to the existence of maximal ideals in all nonzero Boolean rings (the rings satisfying $x^2 = x$ for all x) [14, Definition 2.15, Theorem 4.37] and that is strictly weaker than Zorn's lemma.

In functional analysis, the Hahn-Banach theorem is weaker than Zorn's lemma [24] while the Krein-Milman theorem is equivalent to it [3]. The existence of a Haar measure on a general locally compact Hausdorff topological group was first proved by Weil using Tychonoff's theorem, which makes the existence of Haar measure rely on a property weaker than Zorn's lemma (see the previous paragraph). Cartan later proved the existence of Haar measure using nothing resembling Tychonoff's theorem or Zorn's lemma [9, Chap. 7].

Although Zorn's lemma is used by the overwhelming majority of working mathematicians, there is still a certain degree of caution surrounding it, in the sense that some mathematicians may explicitly be attentive to the use of Zorn's lemma in a proof. Why is this?

The cause is historical and tied up with the equivalence between Zorn's lemma and the Axiom of Choice, since the use of the Axiom of Choice led to extremely counterintuitive theorems in the early 20th century.

- Cantor [7, p. 550] conjectured in 1883 that every set has a well-ordering, calling it a "law of thought" (Denkgesetz). To show $\mathbf{R}$ in particular has a well-ordering was included by Hilbert as part of the first of his famous 23 problems in 1900. Zermelo [29] proved Cantor's conjecture in 1904 by using, for the first time explicitly, the Axiom of Choice.[8] Most mathematicians refused to accept Zermelo's result because a well-ordering on $\mathbf{R}$ (or other uncountable sets) is impossible to imagine and the idea of proving something exists without an explicit procedure to build it was not yet a standard part of mathematics.[9]
- Shortly after Lebesgue introduced measure theory, Vitali [28] in 1905 gave the first example of a non-measurable subset of $[0,1]$ by using a set of representatives for $\mathbf{R}/\mathbf{Q}$ inside $[0,1]$. Measure theory is inspired by basic geometric notions such as length and area, so non-measurable subsets of $\mathbf{R}$ are impossible to imagine. Vitali's construction using a set of representatives for $\mathbf{R}/\mathbf{Q}$ involves the Axiom of Choice.
- In 1905, Hamel [12] proved $\mathbf{R}$ has a basis as a vector space over $\mathbf{Q}$ with the help of Zermelo's Well-Ordering theorem, which had been proved with the Axiom of Choice. Hamel used such a basis to show there are functions $f\colon \mathbf{R} \to \mathbf{R}$ satisfying $f(x+y) = f(x)+f(y)$ that are not of the form $f(x) = cx$ for $c \in \mathbf{R}$. Such functions are highly discontinuous and impossible to imagine in a concrete way.
- In 1924, the Banach–Tarski paradox appeared. It says a closed ball in $\mathbf{R}^3$ of radius $r > 0$ can be broken up into *finitely many* (in fact, five) pieces that can then be rearranged using only *finitely many* rigid motions of space (rotations and translations) to become a closed ball of radius $r' > 0$ no matter what r and r' are, such as a ball the size of a pea being rearrangeable by rigid motions to a ball the size of the Sun. What was counterintuitive about this to mathematicians was *not* closed balls of different radii having the same cardinality, but that the bijection could be achieved using only rigid motions (so no scaling operations). The proof of the Banach–Tarski paradox uses the Axiom of Choice to pick a collection of coset representatives in the group $\mathrm{SO}(3)$ for a certain nonabelian subgroup generated by two elements. Applying all of these coset representatives to a point in the ball will give us a non-measurable subset of the ball.[10] The Banach–Tarski paradox is a refinement of the Hausdorff paradox, which is a peculiar decomposition of a sphere that appeared 10 years earlier and also depends on the Axiom of Choice.

The very negative reaction Zermelo received to his proof of the Well-Ordering theorem led him to propose axioms for set theory in 1908 to clarify his reasoning [30, 31]. This is where the term "Axiom of Choice" (Axiom der Auswahl) first appeared. Few people cared about Zermelo's axioms at first. They became the Zermelo–Fraenkel (ZF) axioms in 1922.

For many years it was unclear if the strange Axiom of Choice (strange due to its paradoxical consequences, as above) was consistent with the rest of set theory. This led to a habit of keeping track of theorems proved with that axiom. In 1938 Goëdel proved that the axioms for ZFC set theory (ZF axioms plus the Axiom of Choice) are consistent if ZF set theory

[8] The Axiom of Choice was used before Zermelo's work without explicit awareness of it as a genuinely new assumption [22, Chap. 1]. It wasn't noticed since earlier consequences were not very controversial.

[9] An example of a non-constructive proof before Zermelo's work was Hilbert's proof in 1890 that all ideals in $\mathbf{C}[X_1,\ldots,X_n]$ are finitely generated. The proof gave no algorithm for finding a generating set.

[10] Since $\mathrm{SO}(3)$ is uncountable and a subgroup with two generators is countable, coset representatives of the subgroup are analogous to representatives for $\mathbf{R}/\mathbf{Q}$ in $[0,1]$ in Vitali's construction of a non-measurable subset of $\mathbf{R}$.

is consistent, so the Axiom of Choice could not lead to a contradiction with ZF unless ZF already has a contradiction within itself. Here is what Joel David Hamkins wrote:[11]

> At least part of the explanation for why people continue to fuss as they do over the Axiom of Choice is surely the historical fact that there was a period of several decades during which the axiom was not known to be relatively consistent with the other axioms of set theory. It was after all not until 1938 that Goedel proved the relatively consistency of ZFC over ZF [...] and several more decades passed until Paul Cohen completed the independence proof by proving that ¬AC is also relatively consistent with ZF [...]. It was during these intermediate times, and especially the time before 1938 when the axiom was not known to be consistent, that the increasingly bizarre consequences of AC were being discovered, and so the habit naturally developed to pay close attention to when the axiom was used. This habit surely lessened after the independence results, but it was not dropped by everyone.

Since Zorn's lemma is equivalent to the Axiom of Choice, the special attention people once paid to the use of the Axiom of Choice in proofs carried over to the use of Zorn's lemma. Although Zorn's Lemma, the Axiom of Choice, and the Well-Ordering theorem are all logically equivalent, their statements don't sound similar at all, so they are not *psychologically* equivalent. Jerry Bona described this state of affairs in the following way, based on his impressions when he was a student:

> The Axiom of Choice is obviously true, the Well-Ordering theorem is obviously false, and who can tell about Zorn's Lemma?

Appendix A. Application to metric spaces

In a real vector space V, the line between vectors v and w is defined to be the set $\{tv + (1-t)w : 0 \leq t \leq 1\}$. A subset of V is called convex if it contains the line between every pair of points in the set. This notion of convexity, while very important in analysis, depends heavily on the real vector space structure: we used real scalars between 0 and 1 and also vector addition. There is a notion of convexity in arbitrary metric spaces, whose definition is based on the idea that the "line" between two points should contain only points in which the triangle inequality is an equality.

In a metric space (M, ρ), a subset S will be called *convex* if for every pair of distinct elements x and y in S there is a $z \neq x, y$ in S such that $\rho(x,y) = \rho(x,z) + \rho(z,y)$. We circumvented the lack of a real vector space structure by not defining the line between x and y, but rather the points that ought to lie on such a "line." (This definition of convex does not quite match the notion in Euclidean space: an *open* star-shaped region of $\mathbf{R}^n$ is not convex in the usual sense but is convex in the abstract sense above. However, for closed subsets of $\mathbf{R}^n$ with the metric induced from $\mathbf{R}^n$, the above notion of convex does match the usual meaning of the term.)

It may appear that our definition is quite weak: we only assume there is one such z. But by repeating the construction with x, y replaced by x, z, we can get more such points, although we don't have much control over the actual distances we can achieve for points "between" x and y. Zorn's lemma will offer that control when we are in a complete space.

[11] See https://mathoverflow.net/questions/22927.

Theorem A.1. *Let (M, ρ) be a complete convex metric space. For distinct points x and y in M and $t \in [0, \rho(x,y)]$, there is a $z \in M$ such that $\rho(x,z) = t$ (and $\rho(x,y) = \rho(x,z)+\rho(z,y)$).*

Proof. We will use Zorn's lemma twice. Also, we will need to use the formulation of completeness in terms of nets, not sequences: every Cauchy net in a complete metric space converges.

First we define some notation. For $a, b \in M$, let

$$[a,b] := \{c \in M : \rho(a,b) = \rho(a,c) + \rho(c,b)\}.$$

For instance, this set contains a and b, and by hypothesis it contains a point besides a and b when $a \neq b$. Intuitively, $[a,b]$ is the set of points lying on geodesics from a to b. It is helpful when reading the following discussion to draw many pictures of line segments with points marked on them. Given t between 0 and $\rho(x,y)$, we will find a $z \in [x,y]$ with $\rho(x,z) = t$.

Some simple properties of these "intervals" are:

(1) $[a,b] = [b,a]$.
(2) If $c \in [a,b]$ and $b \in [a,c]$, then $\rho(b,c) = -\rho(b,c)$, so $b = c$.

Less simple properties are

(3) If $b \in [a,c]$ then $[a,b] \subset [a,c]$ and $[b,c] \subset [a,c]$.
(4) If $b \in [a,d]$ and $c \in [b,d]$ then $[a,c] \subset [a,d]$, $[b,d] \subset [a,d]$, and $[b,c] = [a,c] \cap [b,d] \subset [a,d]$. (We will only need that $[b,c]$ lies in the intersection, not equality.)

Proof of (3): Without loss of generality, we show $[a,b] \subset [a,c]$. For p in $[a,b]$,

$$\begin{aligned}
\rho(a,c) &\leq \rho(a,p) + \rho(p,c) \\
&\leq \rho(a,p) + \rho(p,b) + \rho(b,c) \\
&= \rho(a,b) + \rho(b,c) \\
&= \rho(a,c).
\end{aligned}$$

Therefore $p \in [a,c]$.

Proof of (4): By (3), $[b,d] \subset [a,d]$ and $[b,c] \subset [b,d]$. Therefore $c \in [a,d]$, so $[a,c] \subset [a,d]$, so

$$\begin{aligned}
\rho(a,d) &= \rho(a,c) + \rho(c,d) \\
&\leq \rho(a,b) + \rho(b,c) + \rho(c,d) \\
&= \rho(a,b) + \rho(b,d) \\
&= \rho(a,d).
\end{aligned}$$

Therefore the inequality is an equality, so $b \in [a,c]$, so $[b,c] \subset [a,c]$. Thus $[b,c] \subset [a,c] \cap [b,d]$.

For the reverse inclusion, let $p \in [a,c] \cap [b,d]$. Then

$$\begin{aligned}
\rho(a,d) &= \rho(a,b) + \rho(b,d) \\
&= \rho(a,b) + \rho(b,c) + \rho(c,d) \\
&\leq \rho(a,b) + \rho(b,p) + \rho(p,c) + \rho(c,d) \\
&= \rho(a,b) + \rho(b,d) - \rho(p,d) + \rho(a,c) - \rho(a,p) + \rho(c,d) \\
&= 2\rho(a,d) - \rho(p,d) - \rho(a,p).
\end{aligned}$$

Rearranging terms, $\rho(a,p) + \rho(p,d) \leq \rho(a,d)$, so there is equality throughout, so $\rho(b,c) = \rho(b,p) + \rho(p,c)$. Thus $p \in [b,c]$.

Now we are ready to investigate "geodesics" on M coming out of x. Define a partial ordering on M that might be called "closer to x on geodesics" by

$$z_1 \leq z_2 \text{ if and only if } z_1 \in [x, z_2].$$

In particular, $z_1 \leq z_2$ implies $\rho(x, z_1) \leq \rho(x, z_2)$.

Let's check this is a partial ordering.

If $z_1 \leq z_2$ and $z_2 \leq z_1$, then $z_1 \in [x, z_2]$ and $z_2 \in [x, z_1]$, so $z_1 = z_2$ by (2).

If $z_1 \leq z_2$ and $z_2 \leq z_3$ then $z_1 \in [x, z_2]$ and $z_2 \in [x, z_3]$. By (1), $z_2 \in [z_3, x]$ and $z_1 \in [z_2, x]$. Therefore by (4),

$$z_1 \in [z_1, z_2] \subset [x, z_3],$$

hence $z_1 \leq z_3$.

Define

$$A = \{z \in [x, y] : \rho(x, z) \leq t\}.$$

This set is nonempty, since it contains x. We will apply Zorn's Lemma to A with its induced partial ordering and show a maximal element of A has distance t from x.

Let $\{z_i\}_{i \in I}$ be a totally ordered subset of A. We want an upper bound. Let

$$s = \sup_{i \in I} \rho(x, z_i) \leq t.$$

For each $\varepsilon > 0$, there is some i_0 such that

$$s - \varepsilon \leq \rho(x, z_{i_0}) \leq s,$$

so

$$s - \varepsilon \leq \rho(x, z_i) \leq s$$

for all $i \geq i_0$. For $i_0 \leq i \leq j$, $s - \varepsilon \leq \rho(x, z_i) \leq s$ and

$$s - \varepsilon \leq \rho(x, z_j) = \rho(x, z_i) + \rho(z_i, z_j) \leq s$$

so $\rho(z_i, z_j) \leq \varepsilon$. Thus $\{z_i\}$ is a Cauchy net, so has a limit ℓ by completeness of M. We show this limit is an upper bound in A.

Taking limits,

$$\rho(x, y) = \rho(x, z_i) + \rho(z_i, y) \Rightarrow \rho(x, y) = \rho(x, \ell) + \rho(\ell, y)$$
$$\rho(x, z_i) \leq t \Rightarrow \rho(x, \ell) \leq t.$$

Thus $\ell \in A$.

For $i \leq j$,

$$\rho(x, z_j) = \rho(x, z_i) + \rho(z_i, z_j).$$

Taking limits over j,

$$\rho(x, \ell) = \rho(x, z_i) + \rho(z_i, \ell),$$

so $z_i \in [x, \ell]$, so $z_i \leq \ell$ for all i.

We have justified an application of Zorn's Lemma to A. Let m be a maximal element. That is, $m \in A$, and if $z \in A$ with $m \in [x, z]$ then $z = m$.

Let $B = \{z \in [y, m] : \rho(y, z) \leq \rho(x, y) - t\}$. Since $y \in B$, B is nonempty. Our goal is to show $m \in B$, which is *not* obvious. Note that the definition of B depends on the existence of a maximal element of A.

In B, introduce a partial ordering by $z_1 \leq z_2$ when $z_1 \in [y, z_2]$.

As above, every totally ordered subset of B has an upper bound in B, so by Zorn's lemma B contains a maximal element, m'. Since $m \in [x, y]$ and $m' \in [m, y]$, we get by (4) that

$$[m, m'] \subset [x, m'] \cap [m, y] \subset [x, y].$$

For $z \in [m, m']$,

$$\begin{aligned} z \in [x,y] &\Rightarrow \rho(x,y) = \rho(x,z) + \rho(y,z) \\ &\Rightarrow \rho(x,z) \le t \text{ or } \rho(y,z) \le \rho(x,y) - t \\ &\Rightarrow z \in A \text{ or } z \in B. \end{aligned}$$

Also,

$$\begin{aligned} \rho(x,y) &= \rho(x,z) + \rho(z,y) \\ &\le \rho(x,m) + \rho(m,z) + \rho(z,y) \\ &= \rho(x,m) + \rho(m,y) \text{ since } z \in [m,y] \\ &= \rho(x,y) \text{ since } m \in [x,y]. \end{aligned}$$

Therefore $\rho(x,z) = \rho(x,m) + \rho(m,z)$, so $m \in [x,z]$.

We now have

$$\begin{aligned} \rho(x,y) &= \rho(x,z) + \rho(z,y) \text{ since } z \in [x,y] \\ &\le \rho(x,z) + \rho(z,m') + \rho(m',y) \\ &= \rho(x,m') + \rho(m',y) \text{ since } z \in [x,m'] \\ &= \rho(x,y) \text{ since } m' \in [x,y]. \end{aligned}$$

Therefore $\rho(y,z) = \rho(z,m') + \rho(m',y)$, so $m' \in [y,z]$.

Thus if $z \in A$ then $m \in [x,z] \Rightarrow z = m$. If $z \in B$, then $m' \in [y,z] \Rightarrow z = m'$. Therefore $[m,m'] = \{m,m'\}$, so by *convexity* of M, $m = m'$, hence $m \in A \cap B$. Therefore $\rho(x,m) \le t$ and $\rho(y,m) \le \rho(x,y) - t$, so

$$\rho(x,y) = \rho(x,m) + \rho(m,y) \le \rho(x,y),$$

so $\rho(x,m) = t$. □

The z we constructed in the proof need not be unique. Consider M to be the sphere in $\mathbf{R}^3$ with its surface metric, x and y to be the north and south poles and take z to be a point on some chosen line of latitude. It is natural to expect that if $\rho(x,y)$ is small enough, the z in Theorem A.1 is unique, and this would let us construct *paths* in M. For a proof (without Zorn's lemma!) that every complete convex metric space is in fact path connected, see [6, Theorem 14.1, p. 41].

References

[1] B. Banaschewski, *A new proof that "Krull implies Zorn"*, Math. Logic Quart. **40** (1994), 478–480.

[2] J. L. Bell and F. Jellett, *On the relationship between the Boolean prime ideal theorem and two principles in functional analysis*, Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **19** (1971), 191–194.

[3] J. Bell and D. H. Fremlin, *A geometric form of the axiom of choice*, Fund. Math. **77** (1972), 167–170.

[4] A. Blass, *Injectivity, projectivity, and the axiom of choice*, Trans. Amer. Math. Soc. **255** (1979), 31–59.

[5] A. Blass, *Existence of bases implies the axiom of choice*, pp. 31–33 in: "Axiomatic set theory (Boulder, Colo., 1983)", Amer. Math. Soc., Providence, 1984. Online on http://www.math.lsa.umich.edu/~ablass/bases-AC.pdf.

[6] L. Blumenthal, "Theory and Applications of Distance Geometry," Clarendon Press, Oxford, 1953.

[7] G. Cantor, "Ueber unendliche, lineare Punktmannichfaltigkeiten. 5. Fortsetzung," Math. Ann. **21** (1883), 545–591. Online at https://eudml.org/doc/157080.

[8] I. S. Cohen, *Commutative rings with restricted minimum condition*, Duke Math. J. **17** (1950), 27–42.

[9] J. Diestel and A. Spalsbury, "The Joy of Haar Measure," Amer. Math. Soc., Providence, 2014.

[10] D. Dummit and R. Foote, "Abstract Algebra," 3rd ed., Wiley, New York, 2004.

[11] J. D. Halpern and A. Lévy, *The Boolean prime ideal theorem does not imply the axiom of choice*, pp. 83–134 in: "Axiomatic Set Theory (Proc. Sympos. Pure Math., Vol. XIII, Part I), Amer. Math. Soc., Providence, 1971.

[12] G. Hamel, *Eine Basis aller Zahlen und die unstetigen Lösungen der Funktionalgleichung* $f(x+y) = f(x) + f(y)$ Math. Annalen **60** (1905), 459–462. Online at https://eudml.org/doc/158202.

[13] F. Hausdorff, *Zur Theorie der linearen metrischen Räume*, J. Reine Angew. Math. **167** (1932), 294–311. Online at https://eudml.org/doc/149812.

[14] H. Herrlich, "Axiom of Choice," Springer-Verlag, Berlin, 2006.

[15] W. Hodges, *Krull implies Zorn*, J. London Math. Soc. **19** (1979), 285–287.

[16] P. Jothilingam, *Cohen's theorem and Eakin-Nagata theorem revisited*, Comm. Algebra **28** (2000), 4861–4866.

[17] J. Kelley, *The Tychonoff product theorem implies the axiom of choice*, Fund. Math. **37** (1950), 75–76.

[18] K. Keremedis, *Bases for vector spaces over the two-element field and the axiom of choice*, Proc. Amer. Math. Soc. **124** (1996), 2527–2531.

[19] W. Krull, *Idealtheorie in Ringen ohne Endlichkeitsbedingung*, Math. Annalen **101** (1929), 729–744. Online at https://eudml.org/doc/159366.

[20] T. Y. Lam and M. L. Reyes, *A prime ideal principle in commutative algebra*, J. Algebra **319** (2008), 3006–3027.

[21] S. Lang, "Algebra," 3rd revised ed., Springer, New York, 2002.

[22] G. H. Moore, "Zermelo's Axiom of Choice: its Origins, Development, and Influence," Springer-Verlag, New York, 1982.

[23] A. R. Naghipour, *A Simple Proof of Cohen's Theorem*, Amer. Math. Monthly **112** (2005), 825–826.

[24] D. Pincus, *The strength of the Hahn-Banach theorem*, pp. 203–248 in "Victoria Symposium on Nonstandard Analysis (Univ. Victoria, Victoria, B.C., 1972)," Springer–Verlag, Berlin, 1974.

[25] J. Rotman, "Advanced Modern Algebra," Prentice-Hall, Upper Saddle River, NJ, 2002.

[26] H. Rubin and J. Rubin, "Equivalents of the axiom of choice. II," North-Holland, Amsterdam, 1985.

[27] B. Russell, "The Principles of Mathematics", Cambridge Univ. Press, Cambridge, 1903. Online at https://archive.org/details/principlesofmath01russ.

[28] G. Vitali, *Sul problema della misura dei gruppi di punti di una retta*, Bologna: Tipografia Gamberini e Parmeggiani.

[29] E. Zermelo, *Beweis, daß jede Menge wohlgeordnet werden kann*, Math. Annalen **59** (1904), 514–516. Online at https://eudml.org/doc/158167.

[30] E. Zermelo, *Neuer Beweis für die Möglichkeit einer Wohlordnung*, Math. Annalen **65** (1908), 107–128. Online at https://eudml.org/doc/158340.

[31] E. Zermelo, *Untersuchungen über die Grundlagen der Mengenlehre. I*, Math. Annalen **65** (1908), 261–281. Online at https://eudml.org/doc/158344.

[32] M. Zorn, *A remark on method in transfinite algebra*, Bull. Amer. Math. Soc. **41** (1935), 667–670. Online at https://www.ams.org/journals/bull/1935-41-10/S0002-9904-1935-06166-X/.