Lecture notes for September 22, 2023: Bases and Dimension

version: 10/1/23, 21:43 EDT

In these notes I'll cover most of the content Section 1.6 of Friedberg, Insel, and Spence, *Linear Algebra*, 5th ed. (FIS), minus the examples. However, my order of presentation differs from the book's, and I include some material that's not in Section 1.6.

Recap Proposition 0 below assembles some previously proven facts we'll be using. (For some of these, the proofs may have been so brief that I didn't give them proposition-numbers; some may have been observations, or steps of larger proofs. In case the proofs of parts (a), (b), and (e) below went by too fast in class, or I'm simply imagining that I covered them all this semester, you should find all three of them to be *easy* exercises.)

Proposition 0 *Let V be a vector space and let $S \subseteq V$.*

(a) Let $v \in \text{span}(S)$. Then $\text{span}(S \cup \{v\}) = \text{span}(S)$. [Proven in class.]

(b) If S is linearly independent, then so is every subset of S. [Proven in class.]

(c) Let $S \subseteq V$. Then S is linearly dependent if and ony if there exist $v \in S$ such that $v \in \text{span}(S \backslash \{v\})$ (i.e., iff some element of S is in the span of the other elements). [Proven in class.]

(d) Let $S \subseteq V$ be a linearly independent set, and let v be an element of V that is not in S. Then $S \cup \{v\}$ is linearly independent if and only if $v \notin \text{span}(S)$. (FIS Chapter 1, Theorem 1.7, also proven in class.)

(e) Let A and B be subsets of V.

(i) Let $v \in \text{span}(A \bigcup B)$. Then $v = x + y$ for some $x \in \text{span}(A)$ and $y \in \text{span}(B)$.

(ii) Conversely, any $v \in V$ of the form (element of $\text{span}(A)$) + (element of $\text{span}(B)$) *lies in* $\text{span}(A \cup B)$.

[Proven in class.]

(Equivalent to this bidirectional implication: $\text{span}(A \cup B) = \text{span}(A) + \text{span}(B)$.

[End of recap.]

Basis: definition, existence, and equivalent characterizations

Definition 1 Let V be a vector space. A *basis* of V is a subset $\beta \subseteq V$ that spans V and is linearly independent.

A basis is *finite* if it is a finite set (possibly empty).

Observe that the empty set $\emptyset$ spans the trivial vector space $\{0\}$ and is linearly independent, hence is a basis of this vector space. Conversely, if V is a vector space for which $\emptyset$ spans V, then $V = \{0\}$, and since $\emptyset$ is linearly independent, $\emptyset$ is a basis of V.

Bases (the plural of "basis") can also be characterized as *minimal spanning sets* or as *maximal linearly independent sets*. We define both of these concepts next, and show that each is equivalent to the concept of *basis*.

Definition 2 Let V be a vector space and let $\beta \subseteq V$.

(a) We call β a *minimal spanning set* (for V) if β spans V, but for each $v \in \beta$, the set $\beta \backslash \{v\}$ (β with v removed) does *not* span V.

(b) We call β a *maximal linearly independent set* (in V) if β is linearly independent, but for each $v \in V$ that is not in β, the set $\beta \cup \{v\}$ is linearly dependent.

Note that Definition 2 makes no reference to *cardinality* ("size") of the sets involved. In the usage above, "minimality" (respectively, "maximality") simply means that we lose an indicated property if we remove any element of β (respectively, add any new element to β). Later, we will see how cardinality enters the picture.

Proposition 3 *Let V be a vector space and let $\beta \subseteq V$.*

(a) If β is a minimal spanning set for V, then β is linearly independent.

(b) If β is a maximal linearly independent set, then β spans V.

Proof: (a) Assume that β is a minimal spanning set for V.

Suppose that β is linearly dependent. Let $v \in \beta$ be such that $v \in \text{span}(\beta \setminus \{v\})$. Then from Prop.0(a) (applied with $S = \beta \setminus \{v\}$), it follows that $\text{span}(\beta \backslash \{v\}) = \text{span}(\beta) = V$, contradicting the assumption that β is a *minimal* spanning set for V. Hence β is linearly independent.

(b) Assume that β is a maximal linearly independent set.

Suppose that β does not span V, and let v be an element of V that is not in $\text{span}(\beta)$. Then by Prop. 0(d), $\beta \cup \{v\}$ is linearly independent, contradicting the assumption that β is a *maximal* linearly independent set. Hence β spans V. ■

Corollary 4 *Let V be a vector space and let $\beta \subseteq V$. Then the following are equivalent:*

(i) β is a minimal spanning set for V.

(ii) β is maximal linearly independent set.

(iii) β is a basis of V.

Proof: $\big($(i) $\Longrightarrow$ (iii) and (ii) $\Longrightarrow$ (iii)$\big)$ By Proposition 3, if either (i) or (ii) is satisfied, then β is both linearly independent and a spanning set for V, hence is a basis of V.

$\big($(iii) $\Longrightarrow$ (i) and (ii)$\big)$. Assume that β is a basis of V. Thus, $\text{span}(\beta) = V$ and β is linearly independent.

First suppose that β is not minimal (as a spanning set for V). Let $v \in \beta$ be such that $\text{span}(\beta \backslash \{v\}) = V$. Then, in particular, $v \in \text{span}(\beta \backslash \{v\})$, so by Prop. 0(d), β is linearly dependent, a contradiction. Hence β is minimal (as a spanning set for V).

Next suppose that β is not maximal (as a linearly independent set). Let $v \in V$ be such that $v \notin \beta$ but $\beta \cup \{v\}$ is linearly independent. Then by Prop. 0(c), $v \notin \text{span}(\beta)$, contradicting the assumption that β spans V. Hence β is maximal (as a linearly independent set).

We have now shown that each of (i) and (iii) implies the other, and that each of (ii) and (iii) implies the other. Hence each of the statements (i), (ii), and (iii) implies the other two. ■

Remark 5 In most proofs that three assertions (i), (ii), and (iii) are equivalent, the most efficient way to prove that each assertion is equivalent to every other is to establish a "cycle of implications" such as "(i) $\Longrightarrow$ (ii) $\Longrightarrow$ (iii) $\Longrightarrow$ (i)," (conventional short-hand for "(i) implies (ii), (ii) implies (iii), and (iii) implies (i)"). The short-hand for the logic of the proof above is "(i) $\Longleftrightarrow$ (ii) $\Longleftrightarrow$ (iii)". Although this logical structure requires that four separate implications be shown, rather than just the three needed for a "cycle of implications," for this particular proof the "(i) $\Longleftrightarrow$ (ii) $\Longleftrightarrow$ (iii)" structure seemed more natural to me. The general efficiency-advantage of the "cycle of implications" logical structure becomes more significant when you're showing that four or more statements are equivalent.

Question: Does every vector space *have* a basis?

The answer is yes, but this is not easy to show (and we will not show it). For example, the vector space $\mathcal{F}(\mathbf{R}, \mathbf{R})$ has a basis, but you will never succeed in writing one down! There is one *very* important class of vector spaces for which existence of a basis (in fact, a finite basis) is not hard to establish (and we will do so): the *finite-dimensional* vector spaces, defined below.

Notation 6 (for these notes) For any finite set A, the notation $|A|$ will denote the cardinality of A (a non-negative integer).

Definition 7 A vector space V is *finite-dimensional* if V has a finite spanning set. Otherwise we say that V is *infinite-dimensional.*

Proposition 8 *Let V be a finite-dimensional vector space. Then V has a finite basis.*

Proof: Let $\mathcal{B}$ be the collection of all finite spanning-sets for V. By definition of "finite-dimensional", the set $\mathcal{B}$ is nonempty. Let $\mathcal{N} = \{|A| : A \in \mathcal{B}\}$. Since $\mathcal{B}$ is nonempty, so is $\mathcal{N}$. Thus, $\mathcal{N}$ is a nonempty set of non-negative integers. But every non-empty set of non-negative integers contains a smallest element.[1] Let n be the smallest element of $\mathcal{N}$, and let $\beta \in \mathcal{B}$ be a spanning set for V with n elements.

Suppose that $v \in \beta$ and that $\beta' := \beta \backslash \{v\}$ spans V. Then $\beta' \in \mathcal{B}$, but β' has only $n-1$ elements, contradicting the definition of n. Hence the finite set β is a minimal spanning set for V, so by Cor. 4, β is a basis of V. ■

Remark 9 In view of Prop. 8, a definition of "finite-dimensional vector space" that is equivalent to Definition 7 is: **A vector space V is *finite-dimensional* if (and only if) V has a finite basis.**

Remark 10 In the preceding proof, the basis β we found was not only a minimal spanning set in the sense of Definition 2, it also was minimal in the sense of having smallest *cardinality* among all spanning sets. Our proof established that at least *one* such set β exists, but does not address the question of whether *every* basis of V is minimal in this other cardinality-driven sense. We will address that question later in these notes.

[1] This fact is called the "Well-Ordering Principle", a fundamental property of the integers that is equivalent to the Axiom of Induction and that you may have learned in Sets and Logic. In its usual statement, it asserts that any non-empty set of *positive* integers has a smallest element. However, an easy corollary is that the same holds for any non-empty set of *non-negative integers*.

Finding a basis (given a finite spanning set)

Extracting a basis from a finite spanning set. Suppose that we are *given* a finite spanning set S for a a vector space V. (Possible only if V is finite-dimensional, by Definition 7!) If $S = \emptyset$ or $S = \{0_V\}$, then $V = \{0_V\}$ and $\emptyset$ is a subset of S that is a basis of V, and we're done. If S is nonempty and is not $\{0_V\}$, then we can extract a subset of S that is a basis for V as follows.

> *Informal description of process*: Mentally create "keep", "discard", and "candidate" bins to hold certain elements of S. Initially, the "keep" and "discard" bins are empty, and the "candidate" bin holds all the nonzero elements of S, listed in some order. As we proceed, the sets Cand, Keep, and Disc, consisting of the elements of the "candidate", "keep", and "discard" bins respectively, will change.
>
> Take the first nonzero vector out of the "candidate" bin and put it into the "keep" bin, creating a linearly independent set Keep (initially with just one element). If anything remains in Cand, determine whether the first remaining vector is in span(Keep). If yes, throw that vector into the discard bin. If no, move that vector to the "keep" bin, creating a new, linearly independent Keep set with one more element than before.
>
> Repeat this process until the candidate bin is empty, which must eventually happen since the initial "Cand" set S has only finitely many elements. At each iteration, the new vector added to the linearly independent then-current set Keep is not in the span of what's already there, so the updated "Keep" set is again linearly independent (by Prop. 0(d)). When we're done, every vector we discarded was in the span of previously-kept vectors, so would not have contributed to span(Keep) had we kept it (by Prop. 0(a)). Hence, once the "candidate" bin is empty, the final Keep set (still linearly independent) has the same span as the original set S—namely, all of V—and is therefore a basis of V.

In case there's anything in the above description that's not clear, here is a formal description of the process:

"Reduction Algorithm" (name just for these notes!)

1. Let $S_1 = \{v \in S : v \neq 0_V\} \subseteq S$. (In other words, if S contains the zero vector, remove the zero vector to get S_1. If S does not contain the zero vector, let $S_1 = S$.) Note that any if $0_V \in S$ then for any $v \in \text{span}(S)$ we have

$$v = v' + c0_V = v'$$

for some $v' \in \text{span}(S_1)$. It follows that $\text{span}(S_1) = \text{span}(S)$.

2. If $S_1 = \emptyset$, then $V = \text{span}(S_1) = \{0_V\}$, and $\emptyset \subseteq S$ is a basis of V. In this case, stop; we are done.

3. In this and the remaining steps, we assume $S_1 \neq \emptyset$. Note that, by definition of S_1, all elements of S_1 are nonzero.

   Let $n = |S_1|$ and enumerate the elements of S_1 as $v_1, \ldots, v_n$. Let $T_1 = \{v_1\}$. Since $v_1 \neq 0_V$, the set T_1 is linearly independent. If $n = 1$ then T_1 is a basis of V, and we are done.

   Assume now that $n > 1$. Next, successively check $v_2, \ldots, v_n$ until we either (i) find j such that v_j is not a scalar multiple of v_1 (equivalently, $v_1 \notin \text{span}\{v_1\}$), or (ii) run out of vectors (i.e, we find that v_j is a multiple of v_1 for *every* $j \in \{2, \ldots, n\}$). In case (ii), it is easily seen that $\text{span}(S_1) = \text{span}(T_1)$ (using Prop. 0(a) repeatedly), in which case $T_1 \subseteq S$ is a linearly independent spanning set for V—i.e. a basis of V—and we are done.

   Suppose now that we are in case (i), and let $j_2 \in \{2, \ldots, v_n\}$ be the smallest j for which $v_j \notin \text{span}\{v_1\} = \text{span}(T_1)$ (mentally throwing each successive v_j into a "discard" bin if $v_j \in \text{span}(T_1)$). Let $T_2 = \{v_1, v_{j_2}\}$, a linearly independent set (by Prop. 0(d)) and let $S_2 = T_2 \bigcup \{v_j : j_2 < j \leq n\}$ (i.e. S_2 is S_1 with the "discarded" vectors v_j, $1 < j < j_2$, removed). Since each of the discarded vectors lies in $\text{span}\{v_1\} \subseteq \text{span}(S_2)$, it follows from Proposition 0(a) that $\text{span}(S_1) = \text{span}(S_2)$, hence that $\text{span}(S_2) = V$.

4. Now continue recursively. To have a consistent notational pattern, define $j_1 = 1$. Suppose that $k \geq 2$ and that we have found $j_1, \ldots, j_k$, satisfying $1 = j_1 < j_2 < j_3 < \cdots < j_k \leq n$, (with j_3 present only if $k \geq 3$) such that $T_k := \{v_{j_1}, v_{j_2}, \ldots, v_{j_k}\}$ is linearly independent and $V = \text{span}(S_k)$, where $S_k = T_k \bigcup \{v_j : j_k < j \leq n\}$.

   If $j_k = n$ then $S_k = T_k$, and $T_k \subseteq S$ is both linearly independent and a spanning set for V— i.e. a basis of V—and we are done.

   Assume now that $j_k < n$, and successively check $v_{j_k+1}, \ldots, v_n$ until we either (i) find $j > j_k$ such that $v_j \notin \text{span}(T_k)$, or (ii) run out of vectors (i.e, we find that $v_j \in \text{span}(T_k)$ for *every* $j \in \{j_k + 1, \ldots, n\}$).(*Note*: as a practical matter, for each j this check may require its own algorithm[2], not discussed in this set of notes! However, it is *logically* indisputable that, for each j, either $v_j \in \text{span}(T_k)$ or $v_j \notin \text{span}(T_k)$.) In case (ii), $V = \text{span}(S_k) = \text{span}(T_k)$ (again using Prop. 0(a) repeatedly). Thus, in this case, $T_k \subseteq S$ is a linearly independent spanning set for V, i.e. a basis of V, and we are done.

   Suppose now that we are in case (i), and let $j_{k+1} \in \{j_k+1, \ldots, v_n\}$ be the smallest j for which $v_j \notin \text{span}(T_k)$. Let $T_{k+1} = T_k \cup \{v_{j_{k+1}}\} = \{v_{j_1}, v_{j_2}, \ldots, v_{j_{k+1}}\}$, a linearly

[2]You've already done this type of check in homework exercises. It often involves solving systems of simultaneous linear equations. Eventually, we will have more efficient ways of doing this than we've seen so far.

independent set (by Prop. 0(d)) and let $S_{k+1} = T_{k+1} \bigcup \{v_j : j_{k+1} < j \le n\}$ (i.e. S_{k+1} is S_k with the newly "discarded" vectors—those v_j for which $j_k < j < j_{k+1}$—removed). Since each of the discarded vectors lies in $\text{span}(T_k) \subseteq \text{span}(S_k)$, it follows from Proposition 0(a) that $\text{span}(S_k) = \text{span}(S_{k+1})$, hence that $\text{span}(S_{k+1}) = V$.

Since $j_{k+1} \ge j_k + 1$ and n is finite, eventually this process must terminate. I.e. we will eventually find k for which either $j_k = n$, or for which $j_k < n$ but $v_j \in \text{span}(T_k)$ for every $j \in \{j_k + 1, \ldots, n\}$ ("case (ii)" above). In either of these terminating situations, the set $T_k \subseteq S$ is a basis of V.

Remark 11 Since the span of every subset of a vector space V is itself a vector space, the same algorithm can be used to find a basis of $\text{span}(S)$ when we are given a finite subset $S \subseteq V$.

The description of our algorithm above can easily be rewritten as a proof of the following proposition.

Proposition 12 *Let V be a vector space, and suppose that V has a finite spanning-set S. Then some subset of S is a basis of V.*

Said another way: **every finite spanning-set contains a basis** *(of the ambient vector space).*

■

Note that so far, we still have not addressed the question of whether two bases of a finite-dimensional vector space can have different cardinalities. Nothing we've done (so far) rules out the possibility that, starting with two different finite spanning sets, or with the same finite spanning set ordered in two different ways, our "reduction algorithm" might produce two different bases with different numbers of elements. Using the next theorem below—the "Replacement Theorem"—we'll be able to show that this can't happen, but we're not there yet.

(Note: I am *deliberately* wording this theorem differently from the corresponding theorem in FIS [Theorem 1.10 in Section 1.6], to help you understand the *concepts* by seeing them expressed in more than one way. In accordance with the name chosen for this theorem, my wording introduces notation for the subset of G that's actually being *replaced*, my set R. The set $G \backslash R$ in the wording below is the set H in the book's wording, the subset of G that's *left over* when R is removed from G.)

(Note: The proof of this theorem will be inductive. If you are used to thinking that the terms "base case" and "inductive step" need to appear in an inductive proof, then before reading the proof of the theorem below, see the handout "Inductive proofs: some common mistakes and misconceptions" [particularly, the first full paragraph on p. 2] on the Miscellaneous Handouts page.)

More relations among linearly independent sets, spanning sets, and bases

Theorem 13 ("Replacement Theorem") *Let V be a vector space with a finite spanning set G. If $L \subseteq V$ is a finite linearly independent subset of V, then*

(a) $|L| \leq |G|$ (recall Notation 6), *and*

(b) there exists a subset $R \subseteq G$, having exactly as many elements as L, such that the set $L \cup (G \setminus R)$ spans V. Equivalently, there exists a subset $H \subseteq G$, having exactly $|G| - |L|$ elements,such that $L \cup H$ spans V.

Said yet another way: there is some subset R of G, of the same cardinality as L, such that if we remove the elements of R from G, and replace them with the elements of L, then the new set $G' := L \cup \underbrace{(G \setminus R)}_{H}$ we obtain this way still spans V:

$$V = \text{span}(G) = \text{span}(G').$$

Below, we will give the same proof of this theorem twice: first, a heavily commented version with a diagram that may be helpful (the comments and diagram are not actually part of the proof); second, the same proof with all the comments and diagram removed.

Proof: [The idea behind this proof is essentially to find elements of G, one at a time, that we can replace by elements of L without changing the span; i.e. with the modified version of G still spanning V. The cleanest way to write the argument, however, is to structure it as an inductive proof, proceeding by induction on the cardinality of L, rather than by running one at a time through the elements of a single, fixed (but arbitrary) linearly independent set L. The set R in part (b) is a "discard" set (whose elements we're replacing by elements of L) that's not really important to keep track of; what matters more at the end is the set $H = G \backslash R$ of un-replaced elements.]

We will proceed by induction on the cardinality of L. First suppose that $|L| = 0$. Then $L = \emptyset$, and $H = G$ itself is a subset of G of cardinality $|G| - |L| = |G|$. Trivially, $L \cup H = \emptyset \cup G = G$, so $\text{span}(L \cup H) = \text{span}(G) = V$. Just as trivially, $0 = |L| \leq |G|$ (even if G is empty!). This establishes (a) and (b) in the case $|L| = 0$.

Suppose now that m is a non-negative integer such that statements (a) and (b) are true whenever $L \subseteq V$ is linearly independent set of cardinality $|L| \leq m$. (Above, we showed that 0 is such an m.) Assume now that $|L| = m + 1$ and let $v_1, \ldots, v_{m+1}$ be the elements of L. Let $n = |G|$. Then $L_1 := L \backslash \{v_{m+1}\}$ is an m-element linearly independent subset of V (by Prop. 0(b)), so by the inductive hypothesis, $m \leq n$ ($|L_1| \leq |G|$) and we may select a subset $H_1 \subseteq G$ of cardinality $n - m$ such that $G_1 := L_1 \cup H_1$ spans V. Let $u_1, \ldots, u_{n-m}$ be the elements of H_1.

(This gives us line (1) in the diagram below that runs from line (1) through line (5); the other lines will be explained later. The vertical arrows in the diagram are informal notation, not mathematical arrows; they just indicate certain steps we'll be taking. In line (2), we start examining the (potentially) larger set $G \cup \{v_{m+1}\}$ for a reason that becomes clear only near the end of the proof.

The diagram, like these blue comments, is not part of the proof; I've included it as a guide to make the proof easier to follow.)

$$G_1 = \underbrace{\{v_1, \ldots, v_m\}}_{L_1} \cup \underbrace{\{u_1, u_2, \ldots, u_{n-m}\}}_{H_1} \tag{1}$$

$$\downarrow$$

$$G_1 \cup \{v_{m+1}\} = \underbrace{\{v_1, \ldots, v_m,\ v_{m+1}\}}_{L} \cup \underbrace{\{u_1, u_2, \ldots, u_{n-m}\}}_{H_1\ \text{(still)}} \tag{2}$$

$$\downarrow$$

$$= \{v_1, \ldots, v_m,\ v_{m+1}\} \cup \underbrace{\{u_1, u_2, \ldots, u_{\text{whatever}}, \ldots, u_{n-m}\}}_{H_1\ \text{(still)}} \tag{3}$$

$$\downarrow$$

$$= \{v_1, \ldots, v_m,\ v_{m+1}\} \cup \overbrace{\{\ \underbrace{u_1,}_{\substack{\text{previous} \\ \text{“}u_{\text{whatever}}\text{”}}}\ \ u_2, \ldots, u_{n-m}\}}^{H_1\ \text{(still)}} \tag{4}$$

$$\downarrow$$

$$\big(G_1 \cup \{v_{m+1}\}\big) \backslash \{u_1\} = \underbrace{\{v_1, \ldots, v_m, v_{m+1}\}}_{L\ \text{(still)}} \cup \underbrace{\{u_2, \ldots, u_{n-m}\}}_{H} \tag{5}$$

Since $L_1 \cup H_1$ spans V (part of our inductive hypothesis), in particular $v_{m+1} \in \text{span}(G_1)$, so there exist $x \in \text{span}(L_1)$ and $y \in \text{span}(H_1)$ such that $v_{m+1} = x + y$ (by Prop. 0(e)). Select such x and y. Note that if $y = 0_V$ then $v_{m+1} \in \text{span}(L)$, contradicting the assumed linear independence of L (by Prop. 0(c)). Hence $y \neq 0_V$, implying that $H_1 \neq \emptyset$. Thus $n - m > 0$, so $n - m \geq 1$. Equivalently, $m + 1 \leq n$, which is exactly the statement that $|L| \leq |G|$. This establishes assertion (a) of the theorem for our given set L.

(*Note*: In our diagram, the elements $v_1, \ldots, v_m$ are present only if $m \neq 0$. [This is why our definition of L_1 was "$L_1 = L \backslash \{v_{m+1}\}$" rather than "$L = \{v_1, \ldots, v_m\}$".] Since $n - m \geq 1$, so there is always a u_1 present. The elements $u_2, \ldots, u_{n-m}$ are present only if $n - m \geq 2$. The notation "$H_1\backslash\{u_1\}$" would be more precise notation for "$\{u_2, \ldots, u_{n-m}\}$", being valid even if $n - m = 1$.

Also note that, by definition, all the v's are distinct from each other, and all the u's are distinct from each other. We do not know, or care, whether $v_i = u_j$ for some i, j. However, because it is possible that v_{m+1} is u_j for some j, it is possible that $G_1 \cup \{v_{m+1}\} = G_1$. I.e. the set $G_1 \cup \{v_{m+1}\}$ might not actually have one more element than G_1 does; both sets could be the same. This makes absolutely no difference to our proof; I am simply pointing it out so that you don't think that $G_1 \cup \{v_{m+1}\}$ is *automatically* larger than G_1. That misimpression could confuse you if, for example, you wonder how this proof could possibly handle the case $L = G$.)

Since $y \in \text{span}(H_1)$, there exist scalars $b_1, \ldots, b_{n-m}$ such that $y = b_1u_1 + \cdots + b_{n-m}u_{n-m}$; furthermore, since $y \neq 0_V$, at least one of these scalars must be nonzero. ((1) **The inference that "at least one of these scalars must be nonzero" is <u>crucial</u>; this is what makes the whole argument work!!** (2) Selecting an i for which $b_i \neq 0$, the corresponding element of H_1 is what we've called "u_{whatever}" in line (3) of the diagram.) Without loss of generality we may assume $b_1 \neq 0$. (This amounts to relabeling the elements of H_1 so that "u_{whatever}" in line (3) of the diagram becomes u_1 in line (4).) Let $H = H_1\backslash\{u_1\}$. Then $y = b_1u_1 + y'$ for some $y' \in \text{span}(H)$. (Specifically, $y' = b_2u_2 + \cdots + b_{n-m}u_{n-m}$ if $n - m \geq 2$, and $y' = 0_V$ if $n - m = 1$.) Therefore $v_{m+1} = x + y = x + b_1u_1 + y'$, which implies $b_1u_1 = v_{m+1} - x - y'$, and since $b_1 \neq 0$, it follows that

$$u_1 \quad = \quad \underbrace{\left(\frac{1}{b_1}v_{m+1} - \frac{1}{b_1}x\right)}_{\in \text{ span}(L)} + \underbrace{\frac{-1}{b_1}y'}_{\in \text{ span}(H)} \quad ,$$

the sum of an element of $\text{span}(L)$ and an element of $\text{span}(H)$. Thus $u_1 \in \text{span}(L \cup H)$, implying that $\text{span}(L \cup H \cup \{u_1\}) = \text{span}(L \cup H)$ (by Prop. 0(a)). But $L \cup H \cup \{u_1\} = G_1 \cup \{v_{m+1}\}$. (Compare lines (2) and (5) of the diagram.) Since $\text{span}(G_1)$ is already all of V, so is the span of the more inclusive set $G_1 \bigcup \{v_{m+1}\}$. Hence

$$\text{span}(L \cup H) = \text{span}(L \cup H \cup \{u_1\} = G_1 \cup \{v_{m+1}\} = \text{span}(G_1) = V.$$

(This string of equalities is the reason we "augmented" the set G_1 on line (1) of the diagram to the set $G_1 \cup \{v_{m+1}\}$ on line (2). We needed to show that inserting the element v_{m+1} would allow us to delete one of the elements of H_1 from $G_1 \cup \{v_{m+1}\}$ without "losing span".) This establishes assertion (b) of the theorem for our given set L.

Hence assertions (a) and (b) of the theorem are true whenever $|L| = m + 1$. By induction, these assertions hold for linearly independent subsets $L \subseteq V$ of any finite cardinality. ■

Here is the same proof of the "Replacement Theorem" with all the blue comments and the diagram removed (as well as some simple steps, also in blue previously, that would be okay to treat as obvious to any competent reader without the writer's help).

Proof: We will proceed by induction on the cardinality of L. First suppose that $|L| = 0$. Then $L = \emptyset$, and $H = G$ itself is a subset of G of cardinality $|G| - |L| = |G|$. Trivially, $L \cup H = \emptyset \cup G = G$, so $\text{span}(L \cup H) = \text{span}(G) = V$. Just as trivially, $0 = |L| \leq |G|$. This establishes (a) and (b) in the case $|L| = 0$.

Suppose now that m is a non-negative integer such that statements (a) and (b) are true whenever $L \subseteq V$ is linearly independent set of cardinality $|L| \leq m$. Assume now that $|L| = m + 1$ and let $v_1, \ldots, v_{m+1}$ be the elements of L. Let $n = |G|$. Then $L_1 := L \backslash \{v_{m+1}\}$ is an m-element linearly independent subset of V , so by the inductive hypothesis, $m \leq n$ ($|L_1| \leq |G|$) and we may select a subset $H_1 \subseteq G$ of cardinality $n - m$ such that $G_1 := L_1 \cup H_1$ spans V. Let $u_1, \ldots, u_{n-m}$ be the elements of H_1.

Since $L_1 \cup H_1$ spans V (part of our inductive hypothesis), in particular $v_{m+1} \in \text{span}(G_1)$, so there exist $x \in \text{span}(L_1)$ and $y \in \text{span}(H_1)$ such that $v_{m+1} = x + y$. Select such x and y. Note that if $y = 0_V$ then $v_{m+1} \in \text{span}(L)$, contradicting the assumed linear independence of L. Hence $y \neq 0_V$, implying that $H_1 \neq \emptyset$. Thus $n - m > 0$, so $n - m \geq 1$. Equivalently, $m + 1 \leq n$, which is exactly the statement that $|L| \leq |G|$. This establishes assertion (a) of the theorem for our given set L.

Since $y \in \text{span}(H_1)$, there exist scalars $b_1, \ldots, b_{n-m}$ such that $y = b_1 u_1 + \cdots + b_{n-m} u_{n-m}$; furthermore, since $y \neq 0_V$, at least one of these scalars must be nonzero. we may assume $b_1 \neq 0$. Let $H = H_1 \backslash \{u_1\}$. Then $y = b_1 u_1 + y'$ for some $y' \in \text{span}(H)$. Therefore $v_{m+1} = x + y = x + b_1 u_1 + y'$, and it follows that $u_1 = \left(\frac{1}{b_1} v_{m+1} - \frac{1}{b_1} x\right) + \frac{-1}{b_1} y'$, the sum of an element of $\text{span}(L)$ and an element of $\text{span}(H)$. Thus $u_1 \in \text{span}(L \cup H)$, implying that $\text{span}(L \cup H \cup \{u_1\}) = \text{span}(L \cup H)$. But $L \cup H \cup \{u_1\} = G_1 \cup \{v_{m+1}\}$. Since $\text{span}(G_1)$ is already all of V, so is the span of the more inclusive set $G_1 \bigcup \{v_{m+1}\}$. Hence

$$\text{span}(L \cup H) = \text{span}(L \cup H \cup \{u_1\} = G_1 \cup \{v_{m+1}\} = \text{span}(G_1) = V.$$

This establishes assertion (b) of the theorem for our given set L.

Hence assertions (a) and (b) of the theorem are true whenever $|L| = m + 1$. By induction, these assertions hold for linearly independent subsets $L \subseteq V$ of any finite cardinality. ■

Continuing with new material:

Corollary 14

(a) *In a finite-dimensional vector space, a linearly independent set can never have* <u>*more*</u> *elements than any given, finite, spanning set.* (Thus the "Replacement Theorem" remains true even if we delete the hypothesis that L is finite; that hypothesis is redundant.)

(b) *A finite-dimensional vector space cannot contain an infinite linearly independent set.*

Proof: Let V be a finite-dimensional vector space and let G be a finite spanning set for V. Let $L \subseteq V$ be a linearly independent set.

First suppose that the set L is infinite. Then L has finite subsets of arbitrarily large cardinality. Let $L' \subseteq L$ be a subset of cardinality $|G|+1$. Since any subset of a linearly independent set is linearly independent (Prop. 0(b)), L' is a finite linearly independent set in V whose cardinality exceeds $|G|$, contradicting part (a) of the "Replacement Theorem". Hence L must be finite.

This establishes statement (b). Statement (a) then follows from the Replacement Theorem's part (a). ■

Remark 15 Although the blue comment in Corollary 14 may seem to follow "obviously" from statement (a) of the corollary—technically the comment is premature, because (i) Notation 6 does not tell us what "$|L|$" means if L is an infinite set; (ii) although there is a definition of what *cardinality* means for infinite sets, we have not given that definition in this course; and (iii) even had we given that definition, we'd also have to define what "$\leq$" means when comparing an infinite cardinality (the cardinality of an infinite set) to a finite cardinality (the cardinality of a finite set). However, as we would intuitively expect, it *is* true that when these definitions are given properly, and we use the notation "$|A|$" for the cardinality of an *arbitrary* set, then $|A| \not\leq |B|$ if A is an infinite set and B is a finite set.[3]

Corollary 14 is one of many important corollaries of the "Replacement Theorem". The *most* important of these corollaries is this:

Theorem 16 *Let V be a finite-dimensional vector space. Then all bases of V are finite and have the same number of elements.*

[3] "$|A| = \infty$" is not a mathematically meaningful equation for a general infinite set A. However, *for purposes of reminding ourselves of what Theorem 13(a) would say if we removed the redundant "L is finite" hypothesis*, no great harm is done if we mentally substitute the symbol "∞" for any infinite cardinality, and agree to regard "$\infty > n$" (or "$\infty \not\leq n$") as a true statement for any integer n.

Proof: Let β, β' be bases of V. Since bases are linearly independent sets, and V is finite-dimensional, Corollary 14(b) guarantees us that β and β' are finite sets.

Since β is linearly independent and β' spans V, Theorem 13(a) (the "Replacement Theorem", part (a)) implies that $|\beta| \leq |\beta'|$. Similarly, since β' is linearly independent and β spans V, Theorem 13(a) implies that $|\beta'| \leq |\beta|$. Hence $|\beta'| = |\beta|$. ■

Dimension

Theorem 16 is crucial to the notion of *dimension*:

Definition 17 Let V be a finite-dimensional vector space. The *dimension* of V, written $\dim(V)$, is the cardinality of any basis of V.

Note that this definition makes sense only because of *two* of our previous results: first, that *some* basis of V exists (by Proposition 8); and second, Theorem 16's guarantee that whatever basis of V we choose to count the elements of, we will always get the same number.

Knowing the dimension of a given finite-dimensional vector space V (e.g., knowing that $\dim(\mathbf{R}^n) = n$) can greatly simplifiy the task of determining whether certain subsets of V can possibly be linearly independent or spanning sets:

Proposition 18 *Let V be a finite-dimensional vector space and let $n = \dim(V)$. Then:*

(a) No subset of V with <u>more</u> than n elements can be linearly independent.

(b) No subset of V with <u>fewer</u> than n elements can span V.

(c) For any subset S of V with <u>exactly</u> n elements, the following are equivalent:

- *(i) S is linearly independent;*
- *(ii) S spans V;*
- *(iii) S is a basis of V.*

(Thus, given a vector space V <u>that we **already** know has dimension n</u> (e.g. $\mathbf{R}^n$), and a specific set S of of exactly n vectors in V, if we wish to check whether S is a basis of V it suffices to check <u>either</u> that S is linearly independent <u>or</u> that S spans V; we do not have to check *both* of these properties of a basis.)

Proof: Let β be basis of V (thus $|\beta| = n$) and let $S \subseteq V$.

(a) Assume that $|S| > n$. Then $|S|$ has greater cardinality β, a spanning-set for V. Hence by Theorem 13(a) (the "Replacement Theorem", part (a)), S cannot be linearly independent.

(b) Assume that $|S| < n$. Then the linearly independent set β has greater cardinality than S, so by Theorem 13(a), S cannot be a spanning set for V.

(c) Assume that $|S| = n$. By definition of *basis*, (iii) is equivalent to "(i) and (ii)". Hence it suffices to show "(i) $\iff$ (ii)" (since then if either condition (i) or (ii) holds, they both do).

((i) $\implies$ (ii)) Assume that S is linearly independent. Suppose that S does not span V. Let $v \in V$ be a vector not in $\text{span}(S)$, and let $S' = S \cup \{v\}$. Then S' is linearly independent (by Prop. 0(d)) but has cardinality $n + 1 > n$, contradicting part (a) above.

((ii) $\implies$ (i)) Assume that S spans V. Suppose that S is linearly dependent. Then $S \neq \emptyset$, and there exists $v \in S$ such that $v \in \text{span}(S \setminus \{v\})$ (by Prop. 0(c)). Select such v and let $S' = S \setminus \{v\}$. Then $v \in \text{span}(S')$ and $S' \cup \{v\} = S$, so (using Prop. 0(c)) $\text{span}(S') = \text{span}(S' \cup \{v\}) = \text{span}(S) = V$. But then S' is a spanning set for V with $n - 1$ elements, contradicting part (b) above. ■

Remark 19 Since $\emptyset$ is a basis of the trivial vector space $\{\mathit{0}\}$, it follows that $\dim(\{\mathit{0}\}) = 0$. Conversely, the trivial vector space is the *only* vector space of dimension zero. (If $V \neq \{\mathit{0}\}$ then V contains a nonzero vector, hence a linearly independent set with one element. If V is finite-dimensional, Proposition 18(a) shows that $\dim(V) > 0$. If V is infinite-dimensional, then we currently have no definition of what $\dim(V)$ is, but no harm is done if we simply agree that $\dim(V)$ is *not* zero in this case. If we were digress in order to give a proper definition of $\dim(V)$ for an infinite-dimensional vector space V, we would see that, indeed, $\dim(V)$ is *not* 0 or any other real number.)

Previously we saw that a basis of any vector space is, simultaneously, a maximal linearly independent set and a minimal spanning set (Proposition 4). In a finite-dimensional vector space, Proposition 18 yields the following even stronger statement:

Corollary 20 *Let V be a finite-dimensional vector space and let β be a basis of V. Then*

(a) β has maximal cardinality among all linearly independent subsets of V, and

(b) β has minimal cardinality among all spanning subsets of V.

Proof: Left to student. ■

The student should also think about why the maximality/minimality assertions in Corollary 20 are *stronger* than what Proposition 4, by itself, yields.

Next:

Proposition 21 *Let V be a finite-dimensional vector space and let $S \subseteq V$.*

(a) If S is linearly independent, then S can be "extended" to a basis of V; i.e. there exists some subset $T \subseteq V$, disjoint from S, such that $S \cup T$ is a basis of V.

(b) If S spans V, then S can be "shrunk" to a basis of V; i.e. some subset of S is a basis of V.

Note: (1) Part (b) is exactly Proposition 12; I've restated it here just because it's nicely viewed alongside part (a). (2) In part (a), we did not say that T must be nonempty! It's possible that S is *already* a basis of V, in which case we're going to take $T = \emptyset$.

Proof: (a) Let $n = \dim(V)$ and $m = |S|$. By Proposition 18(a), $m \leq n$. If $m = n$ then Proposition 18(c) implies that $S = S \cup \emptyset$ is already a basis of V.

Now assume $m < n$. By Proposition 18(b), S does not span V. Let $v_1 \in V$, with $v_1 \notin \text{span}(S)$. Then $S_1 := S \cup \{v_1\}$ is linearly independent and has cardinality $m + 1$. If $m + 1 = n$, then, by Proposition 18(c), S_1 is a basis of V. If $m + 1 < n$, then, by the same argument, we may select $v_2 \in V$ such that $S_2 := S_1 \cup \{v_2\} = S \cup \{v_1, v_2\}$ is linearly independent. If $m + 2 = n$ then S_2 is a basis of V. Otherwise, continuing in this fashion, we produce vectors $v_1, \ldots, v_{n-m} \in V$, none of which lies in S, for which $S \cup \{v_1, \ldots, v_{n-m}\}$ is a basis of V.

(b) This is exactly Proposition 12. ■

Dimension of subspaces

Proposition 22 *Let V be a finite dimensional vector space and let $n = \dim(V)$.*

(a) Let $W \subseteq V$ be a subspace. Then:

(i) W is finite-dimensional and $0 \leq \dim(W) \leq n$.

(ii) $\dim(W) = 0$ if and only if $W = \{0_V\}$.

(iii) $\dim(W) = n$ if and only if $W = V$.

(b) Let m be an integer for which $0 \leq m \leq n$. Then V has a subspace of dimension m.

Proof: (a) (i) If $W = \{0_V\}$ then W is finite-dimensional and $\dim(W) = 0$. Assume now that $W \neq \{0_V\}$, and let w_1 be a nonzero vector in W. If $\text{span}\{w_1\} \neq W$, let $w_2 \in W$ be an element not in $\text{span}\{w_1\}$. Then $\{w_1, w_2\}$ is a linearly independent subset of W. Continuing in this fashion, suppose we have produced vectors $w_1, \ldots, w_k \in W$ such that $\{w_1, \ldots, w_k\}$ is linearly independent. If $\text{span}\{w_1, \ldots, w_k\} \neq W$, then the same argument shows that there exists $w_{k+1} \in W$ such that $\{w_1, \ldots, w_{k+1}\}$ is linearly independent. But then Then $\{w_1, \ldots, w_{k+1}\}$ is also a $(k+1)$-element linearly independent subset of V, so by Proposition 18(a), we must have $k + 1 \leq n$; we cannot keep finding vectors in W indefinitely that are not in the span of the already-found $\{w_1, \ldots, w_k\}$. Thus, there must be some $m \leq n$ for which our linearly independent set $\beta_W := \{w_1, \ldots, w_m\}$ is a basis of W. Hence $m \leq n$.

(ii) Shown in Remark 19.

(iii) Clearly if $W = V$ then $\dim(W) = n$. Conversely, suppose that $\dim(W) = n$ and let β_W be a basis of W. Then β_W is an n-element linearly independent subset of V, so by Proposition 18(c), β_W is a basis of V. Hence $W = \text{span}(\beta_W) = V$.

(b) If $n = 0$ then $V = \{0_V\}$, $m = 0$, and V is an m-dimensional subspace of itself.

Suppose now that $n > 0$ and let $\beta = \{v_1, \ldots, v_n\}$ be a basis of V. Let $\beta_m = \{v_i : 1 \leq i \leq m\} = \{v_1, \ldots, v_m\}$. Then β_m is a subset of the linearly independent set β, hence is linearly independent. Let $W = \text{span}(\beta_m)$. Then β_m spans W and is linearly independent, so β_m is a basis of W. Hence W is an m-dimensional subspace of V. ■

Remark 23 If V is a finite-dimensional vector space and the strict inequalities $0 < m < n = \dim(V)$ hold, then V has *infinitely many* subspaces of dimension m. The student should be able to show this easily if $m = 1$, and a little less easily (but not with great difficulty) for larger m.

The student should also be able to show that if V is infinite-dimensional, then V has an m-dimensional subspace for every $m > 0$.

Remark 24 It is important to keep in mind that, given a finite-dimensional vector space V, **the dimension of V is NOT the number of elements of V;** the dimension of V is the *number of elements in a basis of* V. A *basis* of $\mathbf{R}^2$ has more elements than a *basis* of $\mathbf{R}$ (these bases have two elements and one element, respectively), but the vector space $\mathbf{R}^2$ itself does not have "more" elements than the vector space $\mathbf{R}$. In fact, with some cleverness and work, one can write down a bijection $f : \mathbf{R}^2 \to \mathbf{R}$ (a function that is one-to-one and onto; hence it is invertible and the inverse function is, in particular, a surjective [onto] function $f^{-1} : \mathbf{R} \to \mathbf{R}^2$. More generally, for any $m, n > 0$, a bijection from $\mathbf{R}^n$ to $\mathbf{R}^m$ exists.[4] (Exhibiting such bijections is outside the *scope* of this course, though it is not above the *level* of this course.)

[4]Once cardinality of a general, possibly infinite, set is defined, this fact says that $\mathbf{R}^n$ has the same

Unique expansion in terms of a basis

We conclude these notes with the following **extremely** important property of bases. (I've left it for last just to avoid interruping the logical flow of the other material in these notes.)

Proposition 25 ("Unique expansion in terms of a basis")
Let V be a finite-dimensional vector space of dimension $n > 0$ and let $\{v_1, \dots, v_n\}$ be a basis of V. Then for each $v \in V$, there exists a unique n-tuple $(c_1, \dots, c_n)$ of scalars such that

$$v = c_1 v_1 + \cdots + c_n v_n \ . \tag{6}$$

Proof: Let $v \in V$. The basis $\{v_1, \dots, v_n\}$ spans V, so an n-tuple $(c_1, \dots, c_n)$ of scalars exists such that the equality (6) holds. Let $(c_1, \dots, c_n)$ be such an n-tuple, and suppose that $(b_1, \dots, b_n)$ is any n-tuple of scalars for which $v = b_1 v_1 + \cdots + b_n v_n$. Then

$$\begin{array}{rrcl} & b_1 v_1 + \cdots + b_n v_n & = & c_1 v_1 + \cdots + c_n v_n, \\ \Longrightarrow & (b_1 - c_1) v_1 + \cdots + (b_n - c_n) v_n & = & 0_V, \\ \Longrightarrow & b_i - c_i & = & 0, \quad 1 \le i \le n, \end{array}$$

since, being a basis, the set $\{v_1, \dots, v_n\}$ is linearly independent. Hence $(b_1, \dots, b_n) = (c_1, \dots, c_n)$. Thus the n-tuple $(c_1, \dots, c_n)$ for which (6) holds is unique. ■

cardinality as $\mathbf{R}^m$.