Math/CS 120: Intro. to Math | Professor: Padraic Bartlett

Lecture 8: A Crash Course in Linear Algebra

Week 9 | *UCSB 2014*

Qué sed
de saber cuánto!

Pablo Neruda, *Oda a los Números*

1 Linear Algebra

In the past week, we studied the concept of **constructible numbers**, and asked whether numbers like π, $\sqrt[3]{2}$ or $\cos(20^\circ)$ are constructible! In this week's lectures, we will prove that they are not constructible; to do this, though, we need some linear algebra machinery! In these notes we go through the concepts of **vector spaces**, **span**, **basis** and **dimension**, which are the linear algebra concepts we need for our later work.

1.1 Vector spaces, informally.

The two vector spaces you're probably the most used to working with, from either your previous linear algebra classes or even your earliest geometry/precalc classes, are the spaces $\mathbb{R}^2$ and $\mathbb{R}^3$. We briefly restate how these two vector spaces work here:

Definition. The **vector space** $\mathbb{R}^2$ consists of the collection of all pairs (a, b), where a, b are allowed to be any pair of real numbers. For example, $(2, -3), (2, \pi), (-1, 1)$, and $(\sqrt{2}, \sqrt{2})$ are all examples of vectors in $\mathbb{R}^2$. We typically visualize these vectors as arrows in the xy-plane, with the tail of the arrow starting at the origin[1] and the tip of the arrow drawn at the point in the plane with xy-coördinates given by the vector. We draw four such vectors here:

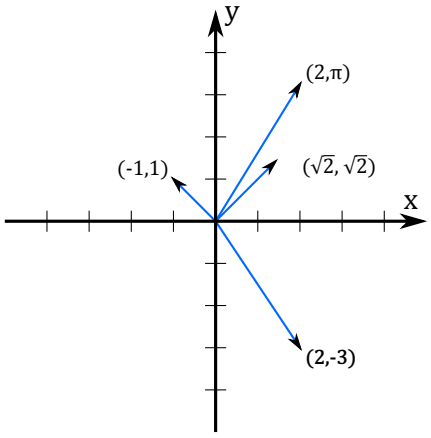

[1] The origin is the point $(0, 0)$ in the plane.

Given a pair of vectors in $\mathbb{R}^2$, we can **add** them together. We do this component-wise, i.e. if we have two vectors (a, b) and (c, d), their sum is the vector $(a+c, b+d)$. For example, the sum of the vectors $(3, -2)$ and $(2, 3)$ is the vector $(5, 1)$.

You can visualize this by taking the arrow corresponding to the first vector that we add, and "translating" this arrow over to the start of the second vector; if you travel along the first vector and then continue along this second translated vector, you arrive at some point in the plane. The arrow connecting the origin to this point is the vector given by the sum of these two vectors! If this seems hard to understand, the diagram below may help some:

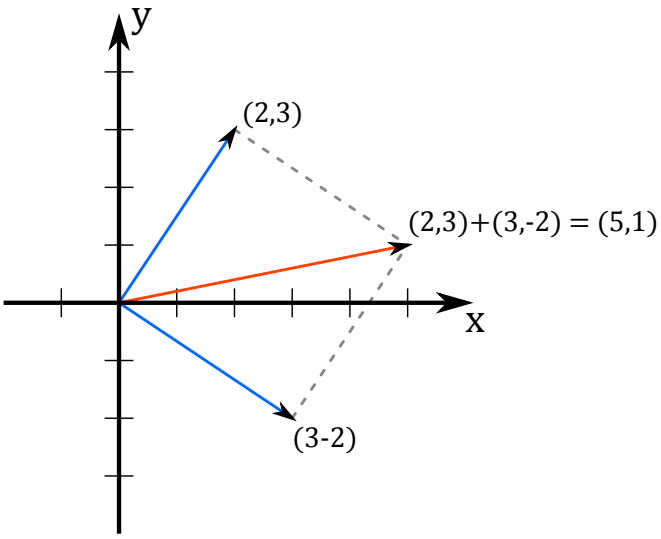

We can also **scale** a vector in $\mathbb{R}^2$ by any real number a. Intuitively, this corresponds to the concept of "stretching:" the vector (x, y) scaled by a, denoted $a(x, y)$, is the quantity (ax, ay). For example, $2(1, 3) = (2, 6)$, and is essentially what happens if we "double" the vector $(1, 3)$. We illustrate this below:

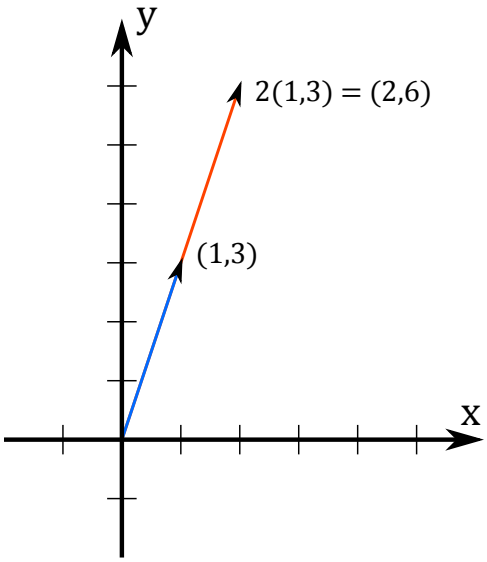

We can define $\mathbb{R}^3$ in a similar fashion:

Definition. The **vector space** $\mathbb{R}^3$ consists of the collection of all pairs (a, b, c), where a, b, c are allowed to be any triple of real numbers. For example, $(0, 1, 2), (3, 0, 2)$, and $(3, 2, 0)$ are all examples of vectors in $\mathbb{R}^3$. We typically visualize these vectors as arrows in three-dimensional xyz-space, with the tail of the arrow starting at the origin and the tip of the arrow drawn at the point in the plane with xyz-coördinates given by the vector. We draw three such vectors here:

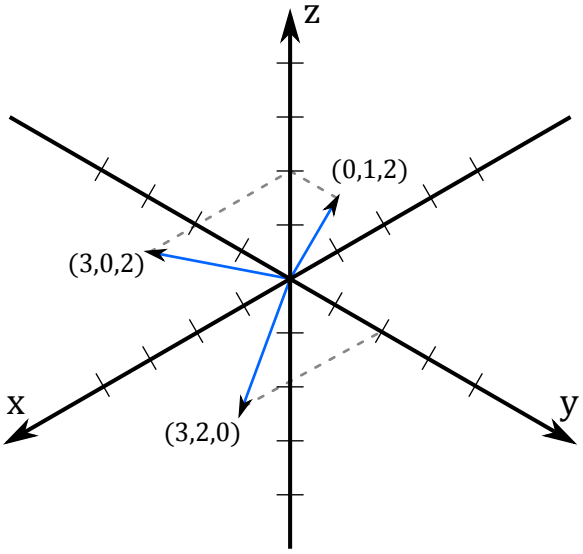

Again, given a pair of vectors in $\mathbb{R}^3$, we can **add** them together. We do this component-wise, i.e. if we have two vectors (a, b, c) and (d, e, f), their sum is the vector $(a+d, b+e, c+f)$. For example, the sum of the vectors $(3, -2, 0)$ and $(2, 1, 2)$ is the vector $(5, -1, 2)$. We can also **scale** a vector in $\mathbb{R}^3$ by any real number a: the vector (x, y, z) scaled by a, denoted $a(x, y, z)$, is the quantity (ax, ay, az). These operations can be visualized in a similar fashion to the pictures we drew for $\mathbb{R}^2$:

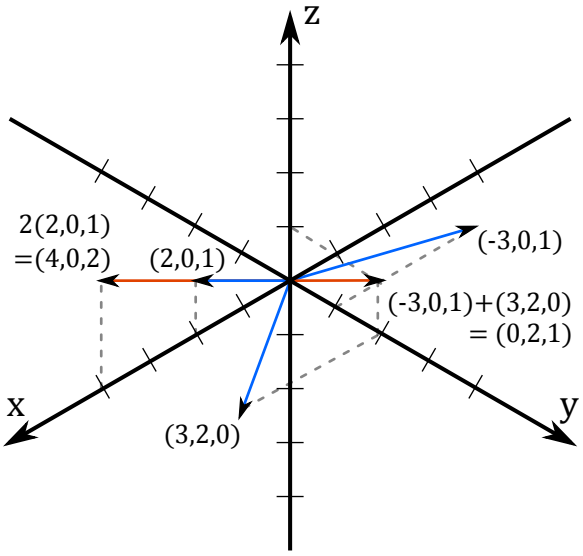

You can generalize this discussion to $\mathbb{R}^n$, the vector space made out of n-tuples of real numbers: i.e. elements of $\mathbb{R}^4$ would be things like $(\pi, 2, 2, 1)$ or $(-1, 2, 1, -1)$.

1.2 Vector spaces, formally.

In general, there are many other kinds of vector spaces — essentially, anything with the two operations "addition" and "scaling" is a vector space, provided that those operations are well-behaved in certain specific ways. Much like we did with $\mathbb{R}$ and the field axioms, we can generate a list of "properties" for a vector space that seem like characteristics that will insure this "well-behaved" nature. We list a collection of such properties and use them to define a vector space here:

Definition. A **vector space** V over a field F is a set V along with the two operations addition and scalar multiplication, such that the following properties hold:

- **Closure(+)**: $\forall \vec{v}, \vec{w} \in V$, we have $v + w \in V$.
- **Identity(+)**: $\exists \vec{0} \in V$ such that $\forall \vec{v} \in V$, $\vec{0} + \vec{v} = \vec{v}$.
- **Commutativity(+)**: $\forall \vec{v}, \vec{w} \in V, \vec{v} + \vec{w} = \vec{w} + \vec{v}$.
- **Associativity(+)**: $\forall \vec{u}, \vec{v}, \vec{w} \in V, (\vec{u} + \vec{v}) + \vec{w} = \vec{u} + (\vec{v} + \vec{w})$.
- **Inverses(+)**: $\forall \vec{v} \in V, \exists$ some $-\vec{v} \in V$ such that $\vec{v} + (-\vec{v}) = 0$.
- **Closure($\cdot$)**: $\forall a \in F, \vec{v} \in V$, we have $a\vec{v} \in V$.
- **Identity($\cdot$)**: $\forall \vec{v} \in V$, we have $1\vec{v} = \vec{v}$.
- **Compatibility($\cdot$)**: $\forall a, b \in F$, we have $a(b\vec{v}) = (a \cdot b)\vec{v}$.
- **Distributivity($+, \cdot$)**: This means two things! On one hand, $\forall a \in F, \vec{v}, \vec{w} \in V$, we have $a(\vec{v} + \vec{w}) = a\vec{v} + a\vec{w}$.

  On the other, $\forall a, b \in F, \vec{v} \in V$, we have $(a + b)\vec{v} = a\vec{v} + b\vec{v}$.

As with fields, there are certainly properties that vector spaces satisfy that are not listed above. For example, consider the following property:

- **New property?(+)**: The additive identity, $\vec{0}$, is unique in any vector space. In other words, there cannot be two distinct vectors that are both the additive identity for a given vector space.

Just like before, this property turns out to be redundant: in other words, this property is implied by the definition of a vector space! We prove this here:

Claim. In any vector space, the additive identity is unique.

Proof. Take any two elements $\vec{0}, \vec{0'}$ that are both additive identities. Then, by definition, we know that because $\vec{0}$ is an additive identity, we have

$$\vec{0'} = \vec{0} + \vec{0'}.$$

Similarly, because $\vec{0'}$ is an additive identity, we have

$$\vec{0} = \vec{0'} + \vec{0}.$$

If we use commutativity to switch the $\vec{0}$ and $\vec{0'}$, we can combine these two equalities to get that

$$\vec{0} = \vec{0'} + \vec{0} = \vec{0} + \vec{0'} = \vec{0'}.$$

Therefore, we have shown that $\vec{0}$ and $\vec{0'}$ are equal. In other words, we've shown that all of the elements that are additive identities are all equal: i.e. that they're all the same element! Therefore, this additive identity element is **unique**: there is no other element that is somehow an additive identity that is different from $\vec{0}$. □

As we did with fields, there are a number of other properties that $\mathbb{R}^n$ possesses that you can prove that any vector space must have: in your textbook, there are proofs that every vector has a unique additive inverse, that $0\vec{v}$ is always $\vec{0}$, that $-1\vec{v} = -\vec{v}$, and other such things.

Instead of focusing on more of these proofs, we shift our attention instead to actually describing some vector spaces!

1.3 Examples of vector spaces.

A few of these are relatively simple to come up with:

- $\mathbb{R}^n$, the example we used to come up with these properties, is a vector space over the field $\mathbb{R}$.
- $\mathbb{C}^n$ is similar. Specifically: $\mathbb{C}^n$ is the set of all n-tuples of complex numbers: i.e.

  $$\mathbb{C}^n = \{(z_1, \dots z_n) | z_1, \dots z_n \in \mathbb{C}\}.$$

  Just like with $\mathbb{R}^n$, we can add these vectors together and scale them by arbitrary complex numbers, while satisfying all of the vector space properties. We leave the details for the reader to check, but this is a vector space over the complex numbers $\mathbb{C}$.
- Similarly, $\mathbb{Q}^n$, the set of all n-tuples of rational numbers

  $$\mathbb{Q}^n = \{(q_1, \dots q_n) | q_1, \dots q_n \in \mathbb{Q}\},$$

  is a vector space over the field $\mathbb{Q}$.
- In general, given any field F, we can form the vector space F^n by taking our set to be

  $$F^n = \{(f_1, \dots f_n) | f_1, \dots f_n \in F\}.$$

  We can add these vectors pairwise: i.e. for any $\vec{f} = (f_1, \dots f_n), \vec{g} = (g_1, \dots g_n) \in F^n$, we can form

  $$(f_1, f_2, \dots f_n) + (g_1, g_2, \dots g_n) = (f_1 + g_1, f_2 + g_2 + \dots f_n + g_n).$$

  We can also scale them: for any $\vec{f} \in F^n, a \in F$, we can form the vector

  $$a(f_1, f_2, \dots f_n) = (a \cdot f_1, a \cdot f_2, \dots a \cdot f_n).$$

  It is not hard to check that because F is a field, F^n is forced to satisfy all of the vector space axioms:

- **Closure(+)**: Immediate. Because F is a field and is closed under addition, the pairwise sums performed in vector addition must create another vector.
- **Identity(+)**: Because F is a field, it has an additive identity, 0. The vector $\vec{0} = (0, 0, \ldots 0)$ is consequently the additive identity for our vector space, as pairwise adding this vector to any other vector does not change any of the other vector's coördinates.
- **Commutativity(+)**: Again, this is a consequence of F being a vector space. Because addition is commutative in F, the pairwise addition in our vector space is commutative.
- **Associativity(+)**: Once more, this is a consequence of F being a vector space. Because addition is associative in F, the pairwise addition in our vector space is associative.
- **Inverses(+)**: Take any $\vec{f} = (f_1, \ldots f_n) \in F^n$. Because F is a field, we know that $(-f_1, \ldots - f_n)$ is a vector in F^n as well. Furthermore, the pairwise addition of these two vectors clearly yields the additive identity $\vec{0}$; therefore, our vector space has inverses.
- **Closure($\cdot$)**: This is a consequence of F being closed under multiplication.
- **Identity($\cdot$)**: Because F is a field, it has a multiplicative identity 1. This 1, when used to scale a vector, does not change that vector at any coördinate because of this multiplicative identity property; therefore 1 is also the scalar multiplicative identity for our vector space.
- **Compatibility($\cdot$)**: This is an immediate consequence from F's multiplication being associative, as for any $a, b \in F$, we have

$$\begin{aligned} a(b(f_1 \ldots f_n) =& a(b \cdot f_1, \ldots b \cdot f_n) = (a \cdot (b \cdot f_1), \ldots a \cdot (b \cdot f_n)) \\ =& (a \cdot b) \cdot f_1, \ldots (a \cdot b) \cdot f_n) = (a \cdot b)(f_1, \ldots f_n). \end{aligned}$$

- **Distributivity($+, \cdot$)**: This is a consequence of F being a vector space. Because multiplication and addition are distributive in F, their combination in our vector space is distributive as well.

- A specific consequence of the above result is that something like $(\mathbb{Z}/5\mathbb{Z})^n$ is a vector space. This is a somewhat strange-looking beast: it's a vector space over a finite-sized field! In particular, it's a vector space with only finitely many elements, which is weird.

  To understand this better, we look at some examples. Consider $(\mathbb{Z}/5\mathbb{Z})^2$. This is the vector space consisting of elements of the form

$$(a, b),$$

  where $a, b \in \{0, 1, 2, 3, 4\}$. We add and scale elements in this vector space using mod-5 modular arithmetic: for example,

$$(2, 3) + (4, 4) = (1, 2),$$

because $2+4 \equiv 1 \mod 5$ and $3+4 \equiv 2 \mod 5$. Similarly,

$$2(3,1) = (1,2),$$

because $2 \cdot 3 \equiv 1 \mod 5$ and $2 \cdot 1 \equiv 2 \mod 5$.

Perhaps surprisingly, these odd-looking vector spaces are some of the most-commonly used spaces in the theoretical computer science/cryptographic settings. In particular, they come up very often in the field of elliptic curve cryptography, and are instrumental to how a number of modern cryptographic schemes work.

There are some odder examples of vector spaces:

- Polynomials! Specifically, let $\mathbb{R}[x]$ denote the collection of all finite-degree polynomials in one variable x with real-valued coefficients. In other words,

$$\mathbb{R}[x] = \{a_0 + a_1 x + \dots a_n x^n | a_0, \dots a_n \in \mathbb{R}, n \in \mathbb{N}\}.$$

  Verifying that this is a vector space is not very difficult — try it if you don't believe it!

- Matrices! Specifically, let $M_{\mathbb{R}}(n,n)$ denote the set of $n \times n$ matrices with real-valued entries. For example

$$M_{\mathbb{R}}(3,3) = \left\{ \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \;\middle|\; a,b,c,d,e,f,g,h,i \in \mathbb{R} \right\}.$$

  If we define matrix addition as simply entrywise addition: i.e.

$$\begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{bmatrix} + \begin{bmatrix} b_{11} & b_{12} & \dots & b_{1n} \\ b_{21} & b_{22} & \dots & b_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ b_{n1} & b_{n2} & \dots & b_{nn} \end{bmatrix} = \begin{bmatrix} a_{11}+b_{11} & a_{12}+b_{12} & \dots & a_{1n}+b_{1n} \\ a_{21}+b_{21} & a_{22}+b_{22} & \dots & a_{2n}+b_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}+b_{n1} & a_{n2}+b_{n2} & \dots & a_{nn}+b_{nn} \end{bmatrix},$$

  and scalar multiplcation as simply entrywise multiplication, i.e.

$$c \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{bmatrix} = \begin{bmatrix} ca_{11} & ca_{12} & \dots & ca_{1n} \\ ca_{21} & ca_{22} & \dots & ca_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ ca_{n1} & ca_{n2} & \dots & ca_{nn} \end{bmatrix},$$

  then this is a vector space! Specifically, it's a vector space for precisely the same reasons that $\mathbb{R}^n$ is a vector space: if you just think of a $n \times n$ matrix as a very oddly-written vector in $\mathbb{R}^{n^2}$, then every argument for why $\mathbb{R}^{n^2}$ is a vector space carries over to $M_{\mathbb{R}}(n,n)$.

  It might seem odd to think of matrices as a vector space, but if you go further in physics or pure mathematics, this is an incredibly useful and common construction.

Two of the most important concepts when working with vector spaces are the ideas of **span** and **linear independence**. We define these in the next subsection:

1.4 Span and linear independence.

When we're working with vectors in a vector space, there's really only two operations we've discussed that we can do: vector addition and scalar multiplication. A useful thing to think about, then, is the collection of all sorts of objects that we can make using these operations! We define these things here:

Definition. Let V be a vector space over some field F. A **linear combination** of some set of vectors S is any sum of the form

$$a_1\vec{v_1} + a_2\vec{v_2} + \dots a_n\vec{v_n},$$

where $a_1, \dots a_n$ are all elements of our field F, and $\vec{v_1}, \dots \vec{v_n} \in S$.

In other words, a linear combination of some set of vectors is anything we can make by scaling and adding these vectors together.

A useful thing to study, given some set $\{\vec{v_1}, \dots \vec{v_n}$ of vectors, is the following: what can we make with these vectors? In other words, what is the collection of all **linear combinations** of these vectors?

This question is sufficiently common that we have a term for its related concept: the idea of **span**!

Definition. Given any nonempty collection of vectors A, the **span** of A, denoted $\text{span}(A)$, is the collection of all linear combinations of elements of A. In other words,

$$\text{span}(A) = \{a_1\vec{v_1} + a_2\vec{v_2} + \dots a_n\vec{v_n} | n \in \mathbb{N}, a_1 \dots a_n \in \mathbb{R}, \vec{v_1}, \dots \vec{v_n} \in A\}.$$

Notice that the span of any nonempty set of vectors $A \subseteq V$ is a vector space in its own right. This is not hard to see: I leave the proof to the HW for interested students.

A genre of questions you'll often encounter in a linear algebra class is "Given some set A of vectors, what is their span? Is some other vector $\vec{w}$ contained within this span?"

We give an example of a calculation here:

Question. You have been hired as a deep-sea navigator. Your first task is to pilot a submarine from our ship, anchored at $(0, 0, 0)$ in the diagram below, to observe a whale at $(12, 6, -6)$. Your submarine has three engines on it. The first, when ran for a cycle, moves the submarine $(2, 1, 0)$ units from its current location. The second, when ran for a cycle, moves the submarine $(0, 2, 1)$ units from its current location. Finally, the third, when ran for a cycle, moves the submarine $(1, 0, 2)$ units from its current location. Submarine engines can be fired for fractional units of time, and when ran in reverse moves the shuttle backwards along that given vector.

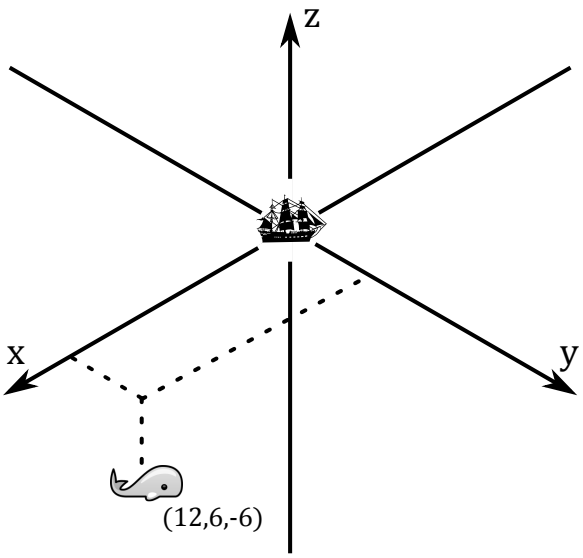

Can you make it to the whale?

Answer. Essentially, our question is asking if the vector $(12, 6, -6)$ is contained in the span of the three vectors $(2, 1, 0), (0, 2, 1), (1, 0, 2)$.

I claim that it is! To see why, simply just start trying to combine these three vectors into $(12, 6, -6)$. If we assume that we fire the first engine for a units, the second for b units, and the third for c units, we're essentially trying to find a, b, c such that

$$a(2, 1, 0) + b(0, 2, 1) + c(1, 0, 2) = (12, 6, -6).$$

This gives us three equations, one for the x-coördinate, one for the y-coördinate, and a third for the z-coördinate:

$$\begin{aligned} 2a + 0b + 1c &= 12, \\ 1a + 2b + 0c &= 6, \\ 0a + 1b + 2c &= -6. \end{aligned}$$

Subtracting two copies of the second equation from the first gives us

$$-4b + c = 0,$$

in other words $c = 4b$. Plugging this into the last equation gives us

$$b + 2(4b) = -6,$$

i.e. $b = -\frac{2}{3}$. This gives us then that $c = -\frac{8}{3}$, and thus that

$$2a - \frac{8}{3} = 12,$$

i.e. $a = \frac{22}{3}$.

Consequently, we've just calculated that

$$\frac{22}{3}(2,1,0) - \frac{2}{3}(0,2,1) - \frac{8}{3}(1,0,2) = (12,6,-6);$$

in other words, that $(12,6,-6)$ is in the span of our set of vectors, and therefore that we can get to it by using our three engines as described by a, b, c!

We can turn span into a verb as follows:

Definition. Suppose that $A = \{\vec{v_1}, \ldots \vec{v_n}\}$ is some set of vectors, drawn from some vector space V. In the proposition above, we proved that span(A) is a subspace of the vector space V: in other words, span(A) is a vector space in its own right! Suppose that span(A) is equal to the vector space U. Then we say that A **spans** U.

This definition motivates a second kind of question: Take some vector space V. Can we find a set A of vectors that spans V?

The answer here is clearly yes: we can just pick A to be V itself! Then the span of A is certainly V, because every vector of V is simply contained in A. Yet, this answer is also kind of... dumb. While A spans V, it does so with a lot of redundancy: i.e. if V was $\mathbb{R}^3$, we'd be using a set with infinitely many elements to span $\mathbb{R}^3$, when we really only need the three vectors $(1,0,0), (0,1,0), (0,0,1)$.

Here's a related question, to help us spot this kind of inefficiency: suppose we have some set A of vectors. Is it possible to remove a vector from A and still have a set with the same span as A?

This concept **also** comes up all the time in mathematics! Therefore, we have a definition for it:

Definition. Let V be some vector space over a field F, and A be a subset of V. We say that the set A is **linearly dependent** if there is some $n > 0$ and distinct elements $\vec{v_1}, \ldots \vec{v_n} \in A$, field coefficients $a_1, \ldots a_n \neq 0 \in F$ such that

$$\vec{0} = a_1\vec{v_1} + a_2\vec{v_2} + \ldots a_n\vec{v_n}.$$

In other words, a set A is linearly dependent if there is a linear combination of elements in A that sums to 0.

If no such combination exists, then we say that A is **linearly independent.**

Notice that if a set is **linearly dependent**, then there is some vector within it that we can remove without changing its span! If this was a linear algebra class, we'd prove this here; instead we reserve it for the HW!

Lemma 1. *Let V be a vector space over a field F, and S be a linearly dependent subset of V. Then there is some vector $\vec{v}$ in S that we can remove from S without changing its span.*

By repeatedly applying Lemma 1, you can get the following theorem:

Theorem 2. *Any finite set of vectors S has a linearly independent subset T, such that span(S) = span(T).*

Side note. In the infinite case, things are much more complicated. In particular, the above process, of repeatedly removing elements one-by-one, isn't guaranteed to ever stop! For example, suppose you had the set $S = \{(z, 0) : z \neq 0 \in \mathbb{Z}\}$.

This is an infinite set. Moreover, because the span of this set is simply $\{(x, 0) : x \in \mathbb{R}\}$, literally any single element $(z, 0)$ of S will have the same span as S: this is because $\frac{x}{z}(z, 0)$ is a linear combination using just $(z, 0)$ that can generate any element in span(S).

However, suppose you were simply removing elements one-by-one from S. It is possible, through sheer bad luck, you would pick the elements $(1, 0), (2, 0), (3, 0), (4, 0), \ldots$ all in a row. This would never get you down to a linearly independent set! In particular, you'd still have all of the elements $(z, 0)$ where z is a negative integer lying around, and there are tons of linear combinations of these elements that combine to $(0, 0)$.

So the argument above doesn't work. However, the result is still true: there **is** a subset that is linearly independent with the same span! Proving that this always exists, though, is tricky, and requires the use of things like the axiom of choice. Talk to me if you want to see how this works!

We define all of this so we can get to the concepts of **basis** and **dimension**:

1.5 Basis and dimension.

Definition. Take a vector space V. A **basis** B for V is a set of vectors B such that B is linearly independent, and span$(B) = V$.

Bases are really useful things. You're already aware of a few bases:

- The set of vectors $\vec{e_1} = (1, 0, 0 \ldots 0), e_2 = (0, 1, 0 \ldots 0), \ldots e_n = (0, 0 \ldots 0, 1)$ is a basis for $\mathbb{R}^n$.
- The set of polynomials $1, x, x^2, x^3, \ldots$ is a basis for $\mathbb{R}[x]$.

As a quick example, we study another interesting basis:

Question. Consider the set of vectors

$$S = \{(1, 1, 1, 1), (1, 1, -1, -1), (1, -1, 1, -1), (1, -1, -1, 1)\}.$$

Show that this is a basis for $\mathbb{R}^4$.

Proof. Take any $(w, x, y, z) \in \mathbb{R}^4$. We want to show that there are always a, b, c, d such that

$$a(1, 1, 1, 1) + b(1, 1, -1, -1) + c(1, -1, 1, -1) + d(1, -1, -1, 1) = (w, x, y, z),$$

and furthermore that if $(w, x, y, z) = (0, 0, 0, 0)$ that this forces a, b, c, d to all be 0. This proves that the span of S is all of $\mathbb{R}^4$ and that S is linearly independent, respectively.

We turn the equation above into four equalities, one for each coördinate in $\mathbb{R}^4$:

$$\begin{aligned} a + b + c + d &= w \\ a + b - c - d &= x \\ a - b + c - d &= y \\ a - b - c + d &= z \end{aligned}$$

Summing all four equations gives us

$$4a = w + x + y + z.$$

Adding the first two equations and subtracing the second two equations gives us

$$4b = w + x - y - z.$$

Adding the first and third, and subtracting the second and fourth gives us

$$4c = w + y - x - z.$$

Finally, adding the first and fourth and subtracting the second and third yields

$$4d = w + z - x - y.$$

So: if $(w, x, y, z) = (0, 0, 0, 0)$, this means that $a = b = c = d = 0$. Therefore, our set is linearly independent.

Furthermore, for any (w, x, y, z), we have that

$$\frac{w+x+y+z}{4}(1,1,1,1) + \frac{w+x-y-z}{4}(1,1,-1,-1)$$
$$+\frac{w+y-x-z}{4}(1,-1,1,-1) + \frac{w+z-x-y}{4}(1,-1,-1,1) = (w,x,y,z).$$

Therefore, we can combine these four elements to get any vector in $\mathbb{R}^4$; i.e. our set spans $\mathbb{R}^4$. □

This example is interesting because its entries satisfy the following two properties:

- Every vector is made up out of entries from ± 1.
- The dot product of any two vectors is 0.

Finding a basis of vectors that can do this is actually an open question. We know that they exist for any $\mathbb{R}^n$ where n is a multiple of 4 up to 664, but no one's found such a basis for $\mathbb{R}^{668}$. Find one for extra credit?

Another natural idea to wonder about is the following: given a vector space V, what is the smallest number of elements we need to make a basis? Can we have two bases with different lengths?

This is answered in the following theorem:

Theorem. Suppose that V is a vector space with two bases $B_1 = \{\vec{v_1}, \dots \vec{v_n}\}$, $B_2 = \{\vec{w_1}, \dots \vec{w_m}\}$ both containing finitely many elements. Then these sets have the same size: i.e. $|B_1| = |B_2|$.

We again reserve this for the HW!

Using this, we can finally define the concept of **dimension**:

Definition. Suppose that V is a vector space with a basis B containing finitely many elements. Then we say that the **dimension** of V is the number of elements in B.

For example, the dimension of $\mathbb{R}^n$ is n, because this vector space is spanned by the vectors $\vec{e_1} = (1, 0, 0 \dots 0), e_2 = (0, 1, 0 \dots 0), \dots e_n = (0, 0 \dots 0, 1)$.