Linear Algebra is the study of linear functions. A function f is said to be **linear** over the field $\mathbb{F}$ (we will use the real numbers ($\mathbb{R}$) for this entire document, but complex numbers ($\mathbb{C}$), rational numbers ($\mathbb{Q}$), and integers modulo a prime p ($\mathbb{Z}_p$) are also common fields) if it satisfies the following two properties:

- $f(cx) = cf(x)$ for all c in the field $\mathbb{F}$, and
- $f(x + y) = f(x) + f(y)$ for all x, y in the domain of f.

So, if $f : \mathbb{R} \to \mathbb{R}$ is a linear function over $\mathbb{R}$, then we see that $f(x) = f(x \cdot 1) = x \cdot f(1)$. Therefore, choosing a value for $f(1)$ determines the entire function, and we see that the set of linear functions is just the set of straight lines through the origin. This discovery should raise a number of questions, such as:

- Isn't $f(x) = 3x + 2$ supposed to be a linear function?
- Why is there an entire branch of mathematics devoted to this set of functions?
- Why is the second property of a linear function even listed if we don't need it?

The answer to the first question varies depending on who you ask and the context. In linear algebra, the answer is a definitive NO. In this setting, $f(x) = 3x + 2$ is called an affine function. To answer the second and third questions, we must broaden our scope somewhat. Instead of having the inputs and outputs of our function be real numbers, we will let them be elements of a vector space.

A **vector space** V over a field $\mathbb{F}$ is a set of objects that can be added together and multiplied by elements of $\mathbb{F}$ to obtain another element of V, along with some other properties needed to make algebra reasonable (if $u, v, w \in V$, $a, b \in \mathbb{F}$, then $v + w = w + v$, $(u + v) + w = u + (v + w)$, $(a + b)v = av + bv$, $a(v + w) = av + aw$, $(ab)v = a(bv)$). Here are some common vector spaces over $\mathbb{R}$:

- $\mathbb{R}$,
- $\mathbb{C}$,
- the set of functions $f : \mathbb{R} \to \mathbb{R}$,
- the set of continuous (smooth, etc.) functions $f : \mathbb{R} \to \mathbb{R}$,
- the set of polynomials with coefficients in $\mathbb{R}$, and
- the set of polynomials of degree $\leq n$ with coefficients in $\mathbb{R}$.

In the case of the last few examples, those who have taken some calculus will recognize the derivative and integral as linear functions over $\mathbb{R}$. There is one more set of common vector spaces over $\mathbb{R}$, the set of n-dimensional vectors with entries in $\mathbb{R}$, denoted by $\mathbb{R}^n$, which I will introduce now.

A **vector** is just an ordered tuple, such as $(1, 4, -3)$. The number of entries in the tuple is the **dimension** of the vector. Vectors are commonly written vertically and with square brackets, like so: $\begin{bmatrix} 1 \\ 4 \\ -3 \end{bmatrix}$. However, vectors and points in n-dimensional Euclidean space are often treated interchangeably, so I will generally use the former notation here to ease typesetting. Vectors of the same dimension can be added merely by adding the entries componentwise, and they can be multiplied by constants (a.k.a. **scalar**s) by multiplying each entry by that constant. So, for example, $(1, 4, -3) + (-5, 2, 4) = (-4, 6, 1)$ and $3 \cdot (1, 4, -3) = (3, 12, -9)$.

When naming a vector, it is common to use a letter with an arrow over it, like $\overrightarrow{v}$. Many textbooks also use bold to indicate a vector, like **v**.

Vectors are often visualized in n-dimensional Euclidean space by drawing an arrow from the origin to the point described by the vector. For example, in $\mathbb{R}^2$, the vector $(2, 1)$ would be drawn as an arrow from the origin to the point $(2, 1)$. Then, the **length** of a vector $\overrightarrow{v}$, denoted $|\overrightarrow{v}|$, is defined in the obvious way. For example, $|(2, 1)| = \sqrt{5}$.

From this point of view, the result of multiplying a vector by a constant c is a vector in the same direction (or exact opposite if $c < 0$) that is $|c|$ times as long. The result of adding two vectors can be found by translating one of them so that it starts at the head of the other, then drawing an arrow from the origin to its head.

Now, we'll consider linear functions whose inputs are vectors. For the sake of example, we'll use $\mathbb{R}^3$ as the domain. Using the properties of linear functions, we see that:

$$\begin{aligned} f((x,y,z)) &= f(x\cdot(1,0,0)+y\cdot(0,1,0)+z\cdot(0,0,1)) \\ &= f(x\cdot(1,0,0))+f(y\cdot(0,1,0))+f(z\cdot(0,0,1)) \\ &= x\cdot f((1,0,0))+y\cdot f((0,1,0))+z\cdot f((0,0,1)). \end{aligned}$$

Therefore, f is completely determined by the choices of $f((1,0,0))$, $f((0,1,0))$, and $f((0,0,1))$. If the outputs of f were 4-dimensional vectors, then we could say $f((1,0,0)) = (2,3,-1,0)$, $f((0,1,0)) = (1,-1,2,-4)$, and $f((0,0,1)) = (-2,3,3,3)$ for example. In this case, the **matrix representation** of f would be:

$$\begin{bmatrix} 2 & 1 & -2 \\ 3 & -1 & 3 \\ -1 & 2 & 3 \\ 0 & -4 & 3 \end{bmatrix}.$$

For clarity, the first column in the matrix is $f((1,0,0))$, the second column is $f((0,1,0))$, and the third column is $f((0,0,1))$. A matrix is a standard way of specifying linear functions whose inputs and outputs are finite dimensional vector spaces. The matrix notation for $f((x,y,z))$ is

$$\begin{bmatrix} 2 & 1 & -2 \\ 3 & -1 & 3 \\ -1 & 2 & 3 \\ 0 & -4 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}.$$

Instead of saying that we evaluate the matrix for a given vector, like we would say using function terminology, we say that we multiply the matrix by the vector. We evaluate the product as follows:

$$\begin{bmatrix} 2 & 1 & -2 \\ 3 & -1 & 3 \\ -1 & 2 & 3 \\ 0 & -4 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = x\cdot\begin{bmatrix} 2 \\ 3 \\ -1 \\ 0 \end{bmatrix} + y\cdot\begin{bmatrix} 1 \\ -1 \\ 2 \\ -4 \end{bmatrix} + z\cdot\begin{bmatrix} -2 \\ 3 \\ 3 \\ 3 \end{bmatrix} = \begin{bmatrix} 2x+y-2z \\ 3x-y+3z \\ -x+2y+3z \\ -4y+3z \end{bmatrix}.$$

It is important to note that it was enough to specify $f((1,0,0))$, $f((0,1,0))$, and $f((0,0,1))$ because every vector in $\mathbb{R}^3$ can be expressed as a sum of multiples (a.k.a. a **linear combination**) of $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$. Furthermore, there were no restrictions on what outputs we could choose because that expression is unique. Thus, we say that the set $\{(1,0,0),(0,1,0),(0,0,1)\}$ forms a **basis** for $\mathbb{R}^3$. This set in particular is known as the **standard basis** for $\mathbb{R}^3$. The elements as listed above are often named $\overrightarrow{e}_1$, $\overrightarrow{e}_2$, and $\overrightarrow{e}_3$, respectively, with a similar notation in other dimensions. In dimension 3, they are also frequently called $\hat{\imath}$, $\hat{\jmath}$, and $\hat{k}$, and it is common to write the vector $(3,2,-1)$ as $3\hat{\imath}+2\hat{\jmath}-\hat{k}$.

Matrices can also be multiplied by other matrices. This is equivalent to function composition. The multiplication is accomplished by multiplying the first matrix by the columns of the second, then adjoining the results as the columns of a new matrix.

Finally, I will give an important example of a linear function $f : \mathbb{R}^2 \to \mathbb{R}^2$: rotation about the origin by an angle θ. It is straightforward to check that this is linear:

- Rotating a vector and then scaling its length gives the same result as first scaling its length and then rotating it.
- Placing two vectors head to tail and then rotating the figure gives the same result as rotating the two vectors and then placing them head to tail.

Then, to determine the function, all we have to do is compute where $(1,0)$ and $(0,1)$ end up after rotating by θ. Basic trigonometry shows that $f((1,0)) = (\cos\theta, \sin\theta)$ and $f((0,1)) = (-\sin\theta, \cos\theta)$. Therefore, the image of (x,y) after rotation is $x\cdot(\cos\theta,\sin\theta) + y\cdot(-\sin\theta,\cos\theta) = (x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta)$. In particular, if we rotate the point $(\cos\phi, \sin\phi)$ by θ, we know by geometry that we get $(\cos(\phi+\theta), \sin(\phi+\theta))$. By our formula, however, we get $(\cos\phi\cos\theta - \sin\phi\sin\theta, \cos\phi\sin\theta + \sin\phi\cos\theta)$. We have just derived the angle sum formulas for sin and cos with seemingly no work at all!

Next time: determinants, dot products, cross products, eigenvalues and eigenvectors, and applications.