Linear Algebra - Review and Application to ODEs

(version 1 of notes)

Linear algebra is the study of vector spaces and linear maps. These ideas go beyond just $\mathbb{R}^n$ and matrices. Let's revisit what they mean and see how they are relevant to the theory of ODEs.

Vectors

Vectors are, roughly, mathematical objects that can be added together and that can be scaled by scalars. Scalars are just numbers; usually we take them to be real numbers but sometimes we allow complex numbers. For addable and scalable things to be called vectors, they need to follow the rules you would expect from addition and scaling, things like

$$\begin{aligned} \vec{v}_1 + \vec{v}_2 &= \vec{v}_2 + \vec{v}_1, \\ c(\vec{v}_1 + \vec{v}_2) &= c\vec{v}_1 + c\vec{v}_2, \end{aligned}$$

and so on. Here vectors are indicated with an arrow like $\vec{v}$, and scalars don't have the arrow, like c. (The arrows are not always there to tell you which are vectors and which are scalars. It's usually clear.) The vectors of a certain type (ones that can be added together) are collected into a set called a *vector space*.

Definition. A *linear combination* of vectors is a sum of scalings of them. So a linear combination of $\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_n$ would be any vector of the form

$$c_1\vec{v}_1 + \cdots + c_n\vec{v}_n$$

for some choice of scalars $c_1, \ldots, c_n$.

Example. Three dimensional space, $\mathbb{R}^3$, is a vector space. The points of $\mathbb{R}^3$, i.e. the vectors, are written with their three coordinates listed vertically. Like this:

$$\begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix}.$$

Below, we see that this vector is a linear combination of $\begin{bmatrix} 1 \\ 4 \\ 6 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$:

$$\begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix} = 1 \begin{bmatrix} 1 \\ 4 \\ 6 \end{bmatrix} + (-2) \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}.$$

Example. Functions are vectors! They can be added and scaled, like in

$$3\sin(x) + 2x^2.$$

Here the function $\sin(x)$ is scaled by three, and added to the function x^2 scaled by two. The result is a new function, one whose value at x is $3\sin(x) + 2x^2$. One says that $3\sin(x) + 2x^2$ is a *linear combination* of $\sin(x)$ and x^2.

Remark. To be clear, the vectors here are *whole functions*, not just their values at a specific choice of x. In other words when I say "the function x^2 is a vector", I'm really talking about a function f which is defined by $f(x) = x^2$. The vector would be f itself, and not any specific $f(x)$.

Remark. (Optional to read this remark...) What would be the vector space in the case of functions as vectors? The main requirement is that vectors from the same vector space need to be addable to produce the same type of vectors. This means that all the functions that get collected together into one "function space" ought to have the same domain. One can require various other things of the functions. One could require them to be continuous on their domain, for example. One could require them to be differentiable, or differentiable with continuous derivative, or twice differentiable with continuous derivative, etc. Here is an example of a function space: $C^7([1, \infty))$. This symbol stands for the vector space of functions defined on the domain $x > 1$ which are seven times differentiable with continous seventh derivative. We will typically not need this level of detail in our treatment!

Linear Independence

Definition. A list of vectors is *linearly independent* if the only linear combination of them that yields zero is the one in which all scalars are chosen to be zero. So the list $\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_n$ is linearly independent if the following happens: The only way to have $c_1\vec{v}_1 + \cdots + c_n\vec{v}_n = \vec{0}$ is to choose $c_1 = c_2 = \cdots = c_n = 0$. A list of vectors is *linearly dependent* if it is not linearly independent.

Example. The vectors $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix}$ are linearly independent. Here's why: if one had

$$c_1 \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + c_2 \begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix},$$

then one would have $3c_1 + 4c_2 = 0$ and $2c_1 + 2c_2 = 0$. It is not hard to conclude from these equations that c_1 and c_2 would both have to be zero.

Example. The functions sin and cos are linearly independent. To see this, suppose you had scalars c_1 and c_2 such that

$$c_1 \sin(x) + c_2 \cos(x) = 0$$

and try to show that c_1 and c_2 must be zero. Remember that this equation holds *for all* x, since the vectors being manipulated are *whole functions*, and the zero on the right hand side is really *the zero function*. So by choosing specific x values we get specific instances of the above equation. Choosing $x = 0$ gives the equation $c_1 0 + c_2 1 = 0$, so $c_2 = 0$. Choosing $x = \pi/2$ gives the equation $c_1 1 + c_2 0 = 0$, so $c_1 = 0$.

Remark. (Optional to read this remark...) You might remember that the *dimension* of a vector space is defined to be the length of the longest possible linearly independent list of vectors. For example, $\mathbb{R}^3$ has dimension 3, because one cannot list any more than three vectors from $\mathbb{R}^3$ while still having the list be linearly independent, and there is a linearly independent list of three vectors: $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$. What about function spaces? What is the dimension of a typical function space? Well, it turns out that there are linearly independent lists of functions that get arbitrarily long, and infinite even. The list of functions $1, x, x^2, x^3, \ldots$ is an example of an infinite linearly independent list. So functions spaces tend to be infinite dimensional.

Linear Maps

Definition. A *linear map* is a vector-valued function of a vector variable that is *linear*. “Linear” means that the function distributes over addition and scaling of vectors. So a linear map is a function $\mathcal{L}(\vec{v})$ of the vector variable $\vec{v}$ with vector output $\mathcal{L}(\vec{v})$ with the property

$$\mathcal{L}(c_1\vec{v}_1 + c_2\vec{v}_2) = c_1\mathcal{L}(\vec{v}_1) + c_2\mathcal{L}(\vec{v}_2).$$

This above property is what is known as “linearity.”

Remark. “Map” is just another word for “function.”

Example. Linear maps with domain $\mathbb{R}^3$ and outputs in $\mathbb{R}^2$ are represented by 2×3 matrices, like

$$\mathcal{L}(\begin{bmatrix} a \\ b \\ c \end{bmatrix}) = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} a + 2b + 3c \\ 4a + 5b + 6c \end{bmatrix}.$$

Example. The derivative operator $\mathcal{L} = \frac{d}{dx}$ is a linear map that operates on functions. It is linear because, as we know, derivatives distribute over sums and scaling:

$$\mathcal{L}(c_1 f(x) + c_2 g(x)) = (c_1 f(x) + c_2 g(x))' = c_1 f'(x) + c_2 g'(x) = c_1\mathcal{L}(f(x)) + c_2\mathcal{L}(g(x)).$$

Actually the same goes for the double derivative operator $\mathcal{L} = \frac{d^2}{dx^2}$ and in fact any number of derivatives $\mathcal{L} = \frac{d^n}{dx^n}$.

Remark. Again, when we treat functions as vectors, the vector is the *whole function.* So in the above example, $\mathcal{L}(f(x))$ can be more accurately written as $\mathcal{L}(f)$, because the whole function f is being fed into the linear map $\mathcal{L}$. This means $\mathcal{L}$ is a function that eats functions!

Example. Here is another random example of a linear map that operates on functions:

$$\mathcal{L}(f(x)) = 5xf''(x) - x^2 f'(x) + 2f(x).$$

Another way to describe what $\mathcal{L}$ does is

$$\mathcal{L} = 5x\frac{d^2}{dx^2} - x^2\frac{d}{dx} + 2.$$

Convince yourself that this $\mathcal{L}$ really is a linear map.

Linear Systems

One of the main things you did in linear algebra was to solve linear systems. This is basically an attempt to "undo" the action of a linear map. The old definition of linear system that you may have been given before is an equation of the form "$A\vec{x} = \vec{b}$," where A is a specific matrix, $\vec{b}$ is a specific vector and $\vec{x}$ is an unknown vector for which you are trying to solve. Let us upgrade to our more sophisticated viewpoint in terms of linear maps:

Definition. A *linear system* is an equation of the form $\mathcal{L}(\vec{x}) = \vec{b}$, where $\mathcal{L}$ is a specific linear map, $\vec{b}$ is a specific vector, and $\vec{x}$ stands for an unknown vector for which one wants to solve.

Now for some terminology that should sound familiar from ODEs:

Definition. A *particular solution* to a linear system $\mathcal{L}(\vec{x}) = \vec{b}$ is a specific choice of $\vec{x}$ that makes it true. The parameterized family of all possible $\vec{x}$'s that are particular solutions is called the *general solution* to the linear system.

Example. Let A be the matrix

$$A = \begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -4 \end{bmatrix}.$$

Let $\mathcal{L}$ be the linear map given by $\mathcal{L}(\vec{x}) = A\vec{x}$. Let $\vec{b}$ be the vector

$$\vec{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

Then $\mathcal{L}(\vec{x}) = \vec{b}$ is an example of a linear system. One particular solution is

$$\vec{x} = \begin{bmatrix} 8 \\ 1 \\ -3 \end{bmatrix}.$$

The general solution is the two-parameter family

$$\vec{x} = \begin{bmatrix} 1 + c_1 - 2c_2 \\ c_1 \\ c_2 \end{bmatrix}.$$

Recall that this is something you could calculate algorithmically, for example by computing a reduced row echelon form.

Example. Let $\mathcal{L}$ be the linear map "$\frac{d}{dx} + 2x$." In other words define $\mathcal{L}$ to be the linear map that operates on functions in the following way:

$$\mathcal{L}(f(x)) = f'(x) + 2xf(x).$$

Let $g(x)$ be the function e^{-x^2}. Then $\mathcal{L}(f(x)) = g(x)$ is an example of a linear system. Here $f(x)$ is in place of the notation $\vec{x}$ and $g(x)$ is in place of the notation $\vec{b}$. The linear system we have written is actually an ODE:

$$f'(x) + 2xf(x) = e^{-x^2}$$

... a linear ODE! We know how to get the general solution by the method of integrating factors. It is the one parameter family

$$f(x) = e^{-x^2}(x + c).$$

General Solution via Homogeneous Linear System

Definition. A *homogeneous linear system* is a linear system of the form $\mathcal{L}(\vec{x}) = \vec{0}$. That is, it is a linear system whose right hand side "$\vec{b}$" is zero.

Suppose you are interested in finding the general solution, i.e. all solutions, of the linear system $\mathcal{L}(\vec{x}) = \vec{b}$. Here $\vec{x}, \vec{b}$ could be $\mathbb{R}^n$-type variables, with $\mathcal{L}$ being represented by a matrix. Or $\vec{x}, \vec{b}$ could be function-type variable with $\mathcal{L}$ being some mish-mash of derivative operations. Either way, there is a way to write down the general solution of $\mathcal{L}(\vec{x}) = \vec{b}$ if you know the general solution of $\mathcal{L}(\vec{x}) = \vec{0}$ and *just one* particular solution of $\mathcal{L}(\vec{x}) = \vec{b}$. Here's how it goes:

1. Find the general solution of the homogeneous version of the linear system: $\mathcal{L}(\vec{x}) = \vec{0}$. Say it turns out to be the n-paramater family

$$\vec{x}_c = c_1\vec{x}_1 + \cdots + c_n\vec{x}_n.$$

2. Somehow find at least one particular solution to the original nonhomogeneous linear system; call it $\vec{x}_p$.

3. Now every solution to the original nonhomogeneous linear system is the sum of $\vec{x}_p$ and a solution to the homogeneous linear system. So the general solution to the nonhomogeneous linear system is the n-parameter family

$$\vec{x} = c_1\vec{x}_1 + \cdots + c_n\vec{x}_n + \vec{x}_p.$$

To understand why this works, we should try to verify the statement in red. First, consider a sum $\vec{x}_p + \vec{x}_c$ where $\vec{x}_p$ is a solution to $\mathcal{L}(\vec{x}) = \vec{b}$ and $\vec{x}_c$ is a solution to $\mathcal{L}(\vec{x}) = \vec{0}$. We have

$$\mathcal{L}(\vec{x}_p + \vec{x}_c) = \mathcal{L}(\vec{x}_p) + \mathcal{L}(\vec{x}_c) = \vec{b} + \vec{0} = \vec{b},$$

so $\vec{x}_p + \vec{x}_c$ is indeed a solution to $\mathcal{L}(\vec{x}) = \vec{b}$. Note the use of linearity there. Now suppose you gave me *any* solution $\vec{x}_1$ to $\mathcal{L}(\vec{x}) = \vec{b}$. Then $\vec{x}_1 - \vec{x}_p$ is a solution to the homogeneous linear system $\mathcal{L}(\vec{x}) = \vec{0}$:

$$\mathcal{L}(\vec{x}_1 - \vec{x}_p) = \mathcal{L}(\vec{x}_1) - \mathcal{L}(\vec{x}_p) = \vec{b} - \vec{b} = \vec{0}.$$

Thus $\vec{x}_1$ really is a sum of $\vec{x}_p$ and a solution to the homogeneous one, namely

$$\vec{x}_1 = \vec{x}_p + (\vec{x}_1 - \vec{x}_p).$$

Example. Recall the linear system

$$\begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -4 \end{bmatrix} \vec{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$$

from above. The general solution to the homogeneous version

$$\begin{bmatrix} 1 & -1 & 2 \\ -2 & 2 & -4 \end{bmatrix} \vec{x} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

is the two-parameter family (plane through the origin)

$$\vec{x}_c = c_1 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} -2 \\ 0 \\ 1 \end{bmatrix}.$$

Since we already had a particular solution $\begin{bmatrix} 8 \\ 1 \\ -3 \end{bmatrix}$ to the original nonhomogeneous linear system, we can deduce that the general solution to the nonhomogeneous linear system is the two-parameter family (shifted plane)

$$\vec{x} = c_1 \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} + c_2 \begin{bmatrix} -2 \\ 0 \\ 1 \end{bmatrix} + \begin{bmatrix} 8 \\ 1 \\ -3 \end{bmatrix}.$$

Example. Recall the linear system (ODE)

$$f'(x) + 2xf(x) = e^{-x^2}$$

from above. The homogeneous version

$$f'(x) + 2xf(x) = 0$$

is separable, and we find that its general solution is the one-parameter family

$$f(x) = ce^{-x^2}.$$

One particular solution to the original nonhomogeneous ODE turns out to be xe^{-x^2}. We can deduce that the general solution to the original nonhomogeneous linear system is the one-parameter family

$$f(x) = ce^{-x^2} + xe^{-x^2}.$$

Remark. In these two examples, we already *had* the general solution to the nonhomogeneous linear system; they were given in previous examples above. So we didn't need to worry about separately solving the homogeneous linear system. But I've done it anyway to demonstrate the principle on something simple. The idea is that we will encounter certain nonhomogeneous linear ODES for which the homogeneous version readily gives us a general solution, and for which we will struggle to pry even one solution out of the fully nonhomogeneous version. Then the principle stated here will allow us to write down a general solution to the nonhomogeneous linear ODE.