

When you're thinking about how to visualize something challenging, it's worth first thinking about what information you need to represent, and then what ways you might represent that information.

We live in three spatial dimensions and one time dimension, which gives us four independent possible representations for information (3 spatial + 1 time). However, because monitors and paper are two-dimensional, this number is reduced to three (2 spatial + 1 time) until we have reliable holographic imaging (or unless you use 3D printing). In addition, we have three types of cones, which gives us three dimensions of color information. However, because about 10% of color people have only two of these three types of cones, it is not recommended to use color to represent more than one dimension of information (except when representing color itself).

We have considerable capacity for obtaining information via sound, but except for time/frequency information (and specifically, representing sound) it's not a great way to convey arbitrary scientific data. Smell, taste and tactile data are generally worse.

So we have generally just four dimensions to convey information: 2 space + 1 time + 1 color.

0.1 Features of dimensions

We have these four dimensions to play with, and each of them have strengths and weaknesses.

0.1.1 Spatial dimensions

The two spatial dimensions are the workhorses of scientific visualization. In the classic plot, each represents a separate dimension of data (e.g. $f(t)$ vs. t). The spatial dimensions excell at representing data quantitatively, since we can carefully measure the position on a monitor or paper.

0.1.2 Temporal dimension

Time is a tricky dimension to use in representing data, and is most frequently used to represent time itself, i.e animation. This has the advantage that the human brain is excellent at processing time-dependent data.

Animation has the major disadvantage that the viewer cannot simultaneously view representations shown at different times, and instead must remember what they already saw, and wait for what they will later see. Furthermore, it is extremely difficult to extract data from the time-dependence of an animation, so while animations may give us a sense of what something is doing, it is hard for instance to extract actual rates of change.

0.1.3 Color dimensions

The color dimension requires care both due to color blindness, but also because colors are not reliably reproduced across devices.

One use for color is to represent a discrete set of data (i.e. which curve is this one). This works well, but should not be the sole distinguishing feature in publication, because papers may be printed in grayscale, or readers may be color blind (although default sets of discrete colors are usually designed to be distinguished by red-green color blind viewers).

For communication of a continuous variable, even grayscale is not reliable (see gamma correction). As a result, if color is to be used quantitatively, a color bar is essential. Moreover, a perceptually uniform color scheme is close to essential.

0.2 Visualising two continuous coordinates

If you have only two coordinates to plot, use the same kind of plot you learned in K-12. This makes both dimensions readily measurable. It's a classic plot because it's effective.

0.3 Visualising three continuous coordinates

The rest of this section will be about what to do when you have *three* coordinates to plot, which could for instance be x, y, z , or $x, \psi(x, t), t$, or $V(x, y)$, x , y . The ways to represent such a data set include either representing each coordinate using a distinct visualization dimension (e.g. heat maps or animated plots), using two visualization dimensions to approximately represent three data coordinates (e.g. 3D plots), or discretizing one of the data coordinates and displaying only a subset of that information (e.g. contour plots, or translated plots).

0.3.1 3D plots

The most popular representation among students is a “3D” plot. This probably feels natural when you have three coordinates to plot, because it has 3 in its name. This family of representations use the two spatial dimensions to attempt to convey three coordinate dimensions. This is an inherently lossy process. At the most fundamental, the same x, y point on the page necessarily corresponds to an entire line of coordinate space, which introduces one coordinate of ambiguity. In addition, if there happen to be more than one datum along this line of sight, information will be blocked. These limitations make getting data *out* of a 3D plot essentially impossible. A representation where you can put your data in, but cannot take it back out is okay for amusement and communication of ideas to a broader audience, but 3D plots are severely limited as a means of scientific communication.

In the case where none of the three coordinates is a function of the other two, a 3D plot may be your only choice. An example would be the visualization of a 3D protein structure. In addition, when the three dimensions to be viewed are themselves spatial dimensions, a 3D visualization may be the most effective.

There are a few ways to somewhat compensate for the limitations of a 3D plot. The most common one is to animate the plot by rotating it. This can resolve the ambiguity regarding where a given object is located by changing the perspective, and it can bring into view objects that were blocked from a different perspective. However, it makes it all the more impossible to extract any data from the plot, since measurement becomes impossible.

Another popular mitigation of this approach—particularly when one coordinate is not spatial—is to redundantly represent that coordinate using color. This certainly helps to reduce the ambiguity in the representation, but ultimately this just creates a difficult-to-read heat map.

0.3.2 Animated plots

Animated plots can be quite effective when one of your coordinates is time. They suffer, however, the same limitations discussed above for any animation, specifically in that it is very hard to extract data from an animation.

0.3.3 Heat maps

A heat map can convey three dimensions of information, provided one dimension is a function of the other two. Heat maps tend to be very effective.

0.3.4 Translated plots

A nice trick can be to discretize one of the coordinates, and then represent it using an ordinary plot with translation. The effect is similar to a 3D plot from a particular perspective, and like a 3D plot it uses the two spatial dimensions to represent all 3 dimensions. However, a translated plot can still enable data to be quantitatively extracted.

0.3.5 Contour plots

Contour plots are the black-and-white siblings of heat maps, and like heat maps are limited to the case where one coordinate is a function of the other two. They discretize the dependent coordinate is represented by drawing curves in space. The contour plot does allow data to be quantitatively extracted, but the discretization of the dependent coordinate omits considerable information.

0.3.6 Slice plots

“Slice plots” (my name) look very much like a contour plot, but the coordinate that is discretized is not the dependent coordinate, but instead one of the other two. Multiple curves are drawn representing the discretized coordinate values. Slice plots are most effective when there is a monotonic relationship that prevents slices from overlapping or if all overlaps occur at the same point. The slice plot does allow data to be quantitatively extracted, but the discretization of one independent coordinate omits considerable information (but different information from the contour plot).

0.4 Visualising four continuous coordinates

With four continuous dimensions (e.g. $\Phi(x, y, z)$, x, y, z , or $\psi(x, y, t)$, x, y, t) we start hitting the limits of what we can visualize. I could identify just one way to represent an entire four dimensions effectively, which I wouldn't strongly recommend unless one of the dimensions is time. All other methods require visualizing only a subset of your data.

Most of these approaches involve picking a subspace of your data to plot, and then using one of the approaches above to plot fewer-dimensional data.

0.4.1 Animated heat maps

You can actually take any static visualization that can represent three coordinates and animate it to obtain an animated visualization of four coordinates. So long as the animated coordinate is time, this can be effective.

0.4.2 3D contour surfaces

You can plot a contour surface, provided one coordinate is dependent on the other three, as is most often the case. This suffers from all the limitations of 3D plots above, on top of the factor of showing only a few values for the dependent coordinate. Blocking of view can be mitigated by making the surfaces partially transparent. As with any 3D plot, you can get very little data out of a contour surface plot. Basically, you can extract the existence of regions with the specified contour value.

0.4.3 2D slice heat maps

We can choose a plane in 3 dimensions to reduce the total number of coordinates from four down to three, and then use a heat map.

0.4.4 2D slice contour plots

We can choose a plane in 3 dimensions to reduce the total number of coordinates from four down to three, and then use a contour plot.

0.4.5 2D slice 3D plots

I only mention this method because it is extremely popular among students. It is a terrible representation! We could choose a plane in 3 dimensions to reduce the total number of coordinates from four down to three, and then use a 3D plot to represent those three dimensions. This has all the problems that 3D plots have, but it's got those problems gratuitously by attempting to represent a non-spatial coordinate (often $\Phi(x, y, z_0)$) using one of the two spatial dimensions that are already occupied representing the actual spatial dimensions.

0.4.6 1D line plots