MONTAGE

Information Visualisation - The Science Underlying Today's Best Practices

Chip Felton – Principal Consultant

HEAD OFFICE
Level 3, 61 Cambridge Tce, Christchurch 8013
POSTAL
PO Box 8953, Riccarton, Christchurch 8440

FREEPHONE ▪ 0800 666 824
PHONE ▪ +64 3 962 6050
www.montage.co.nz

Information Visualisation - The Science Underlying Today's Best Practices

Charts and other graphics are a powerful means for sharing information derived from data. But how do you decide what chart or other graphical representation to use to best convey that information? Are such decisions just a matter of personal preference, taste or aesthetics? Or is there a firmer, more scientific basis for choosing the best visual representation?

For instance, let's say you are looking at sales data and you want to compare product categories based on what percentage of total sales each category represents. What is the best way to communicate the results of that comparison visually? Should you choose a bar chart, a pie chart, or perhaps just use a data table? Or let's take a more complex example: say you are designing a dashboard for use by the company executive in charge of sales. Among the many types of visual representations of sales information such a dashboard could contain, such as point in time summaries, sales trends over time, best and worst performing products or stores or sales representatives, what visual design choices should you make to ensure that the dashboard always communicates the most important pieces of information at a glance?

In each case, the visual representations you choose should be ones that communicate the data's information content quickly and accurately. But which visual representations possess these desired attributes and which do not? Is there a body of knowledge and practice out there that can guide BI practitioners, data analysts, and other knowledge workers to choose the right visualisation every time?

It turns out there is. The field is called Information Visualisation, and it is based on a growing body of evidence derived from the scientific study of human visual perception. In a nutshell, Information Visualisation aims to derive best practice for the visual representation of data from what has been learned about how human vision works; in particular, how we go about recognising visual patterns in the world around us.

Over the past several decades, researchers from various disciplines have built up a detailed understanding of how it is we - quickly and largely without conscious effort - make sense of the sensory data constantly flooding in through our eyes. It turns out that evolution has optimised human vision to pick out particular features within all that incoming data. Some of these features relate to colour, some to spatial position, others to the form of objects, and others to motion.

For instance, concerning colour, our vision is optimised to rapidly detect differences in hue and intensity. To experience the power of this directly, let's take on a visual task: how many '5's are contained in each of the two tables below?

Table 1

47502750326304746520105846295723741872373 4
92456123024567235766522312128565753236263 5
47230564820657349466251566808782544472334 4
68375684461234567251215467563521332546712 3

Table 2

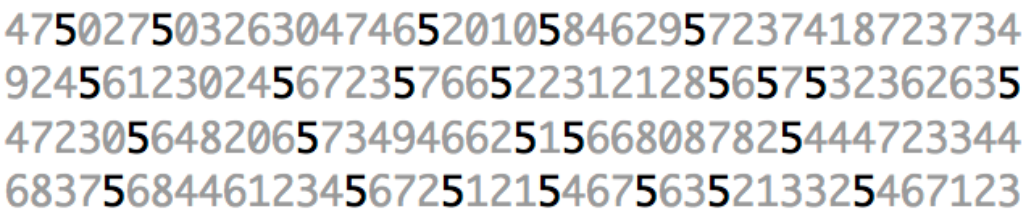

To find the answer using Table 1, you have to read across each line, examining each number individually to determine if it is a '5'. In Table 2, each '5' has been highlighted by increasing its colour intensity. As a result, you don't have to read the whole table and you can quickly count up the number of '5's. Thus, by leveraging the fact that our vision has been optimised by evolution to quickly discern differences in colour intensity, Table 2 does a much better job of visually communicating the important bits of information more quickly and clearly.

The Information Visualisation take-away here is that simply highlighting what is most important in your data display is a powerful aid to quick, accurate comprehension. For instance, consider the two versions of a simple Summary Statistics table displayed below from a hypothetical sales dashboard. In the first version, all the numbers in the table are treated equally visually. However, the second version leverages our ability to rapidly detect differences in colour intensity and hue to draw the eye to one key statistic: a missed sales target. While this key fact also appears in the first version, the viewer has to study the entire table to find it. For a busy executive, this one visual design feature might mean the difference between noticing or overlooking a key fact about the business's current performance that might require action.

Summary Stats Table - Version 1

Sales Summary Stats for 2013 Q4

Region	Department	Actual	+/- Target
Central	Furniture	$85.6K	$4.3K
	Office Supplies	$117.6K	$5.9K
	Technology	$136.4K	$6.8K
East	Furniture	$119.6K	$6.0K
	Office Supplies	$66.2K	$3.3K
	Technology	$151.1K	$7.6K
International	Furniture	$842.7K	$42.1K
	Office Supplies	$868.5K	$43.4K
	Technology	$1,203.4K	$60.2K
South	Furniture	$32.7K	$1.6K
	Office Supplies	$31.9K	$1.6K
	Technology	$119.3K	-$17.9K
West	Furniture	$78.3K	$3.9K
	Office Supplies	$47.3K	$2.4K
	Technology	$88.9K	$4.4K

Summary Stats Table – Version 2

Sales 2 for 2013 Q4

Region	Department	Actual	+/- Target
Central	Furniture	$85.6K	$4.3K
	Office Supplies	$117.6K	$5.9K
	Technology	$136.4K	$6.8K
East	Furniture	$119.6K	$6.0K
	Office Supplies	$66.2K	$3.3K
	Technology	$151.1K	$7.6K
International	Furniture	$842.7K	$42.1K
	Office Supplies	$868.5K	$43.4K
	Technology	$1,203.4K	$60.2K
South	Furniture	$32.7K	$1.6K
	Office Supplies	$31.9K	$1.6K
	Technology	$119.3K	-$17.9K
West	Furniture	$78.3K	$3.9K
	Office Supplies	$47.3K	$2.4K
	Technology	$88.9K	$4.4K

We have so far focussed on the impact of leveraging just one key pattern recognition feature of human vision. All the other features can be similarly leveraged to produce optimal information displays. Without going into more detail at this time, here are some of the many practical highlights:

Concerning form, it turns out that evolution has made us much more accurate at comparing line lengths than angle sizes. This means that, when we need a chart to communicate part to whole relationships such as percentage of total sales by product category, the relative contributions are more accurately conveyed using a bar chart than a pie chart because humans are better at discerning differences in length than they are in comparing slices of a circle.

Here's an example (below) of a bar chart representation of sales by product category versus a pie chart representation of the same data. While the pie chart does tell you which four categories account for the most sales (office machines, tables, telephones, chairs), the bar chart conveys the same information and also makes it possible to quickly compare those four categories to each other accurately, something you can't do with the pie chart without adding explicit labelling.

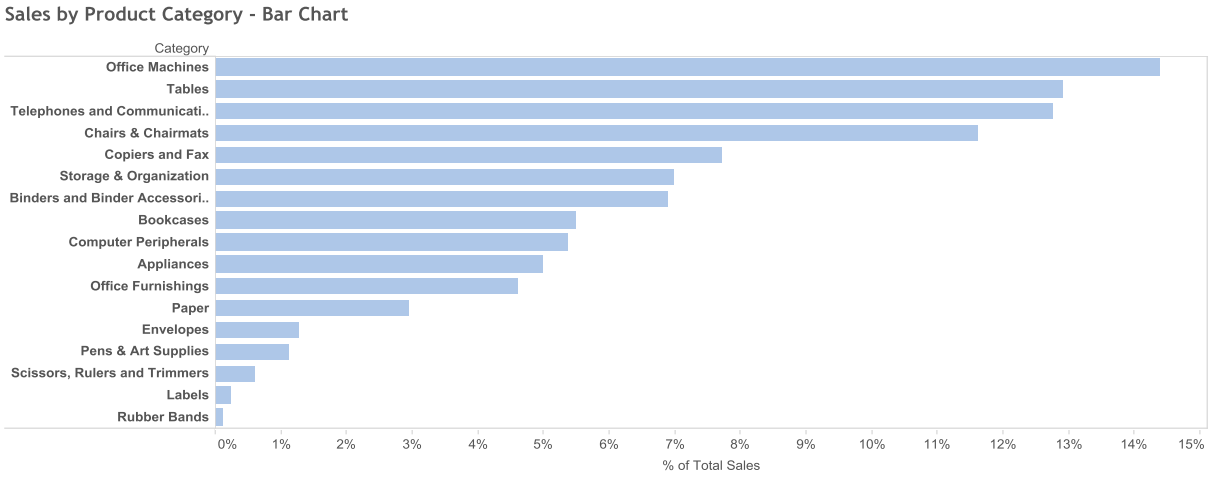

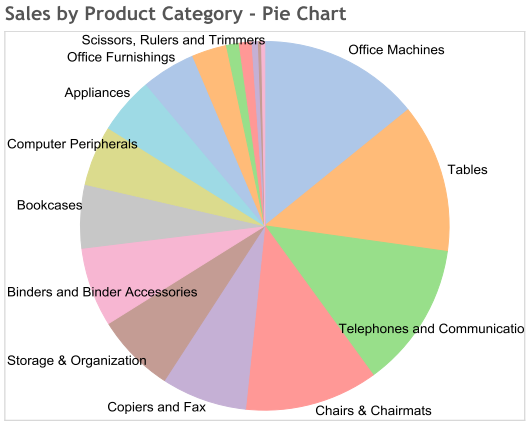

We are also good at detecting differences in the 2-D position of objects. This, combined with our facility for comparing line lengths, means that line charts work well for communicating such things as change over time (trends). Note though that so-called 3-D versions of 2-D chart types are generally not good choices. When displayed on a two

dimensional surface such as a computer screen or paper, 3-D charts use various projection techniques to provide the illusion of 3-D; in turn, these projections can introduce substantial distortion and undermine our ability to interpret length and 2-D position differences accurately (3-D graphics nevertheless are an important tool in scientific visualisation, where the goal is to simulate how real-world objects appear in space).

Concerning colour hue, it's best to use bright fully saturated colours sparingly in information visualisations. These are best used to highlight the most important facts or differences, with the bulk of the visualisation painted using more subdued colours commonly found in nature. While judicious use of bright colour can effectively focus the viewer's attention, too much distracts and thus undermines clear, rapid communication.

Resources to learn more

If you want to learn more about Information Visualisation (and I urge anyone involved in producing information graphics to do so), the best place to turn initially are the books written by Stephen Few (see "References for further reading" below). This brief article owes much to his work. All BI practitioners owe Stephen Few a debt in that he has virtually singlehandedly brought together in accessible form the practical insights and best practices found in more technical works on Information Visualisation. He has also relentlessly hounded the major BI product vendors to improve their visualisation tools so that they support Information Visualisation best practices and make it easier for BI practitioners to apply best practices to their work.

In addition to Stephen Few's books, his website Perceptual Edge (http://www.perceptualedge.com) contains a wealth of practical resources. Overall, if you internalise the lessons from these books, you'll be well equipped to create effective information graphics.

For those interested in diving more deeply into the scientific aspects of Information Visualisation, I would recommend the book "Information Visualisation: Perception for Design" by Colin Ware. There is also the book "Readings in Information Visualization: Using Vision to Think" that organises together a set of seminal papers that helped to define the field of Information Visualisation.

References for further reading

1. Information Dashboard Design: The Effective Visual Communication of Data, Stephen Few, O'Reilly Media, 2006
2. Now You See It: Simple Visualization Techniques for Quantitative Analysis, Stephen Few, Analytics Press, 2009
3. Show Me the Numbers: Designing Tables and Graphs to Enlighten, Second Edition, Stephen Few, Analytics Press, 2012
4. Information Visualization, Second Edition: Perception for Design , Colin Ware, Morgan Kaufmann, 2004
5. Readings in Information Visualization: Using Vision to Think Stuart K. Card (Editor), Jock Mackinlay (Editor), Ben Shneiderman (Editor), Morgan Kaufmann, 1999