

Data Visualisation


Richard Wright, April 2006

In: Matthew Fuller (ed), "Software Studies", MIT Press, 2007.

Any transformation of digital material into a visual rendering can arguably be called a visualisation, even the typographic treatment of text in a terminal window. Conventionally however, "data visualisation" is understood as a mapping of digital data into a visual image. The need for visualisation was first recognized in the sciences during the late eighties as the increasing power of computing and the decreasing cost of digital storage created a surge in the amount and complexity of data needing to be managed, processed and understood. In 1987 the US National Science Foundation published their "Visualisation in Scientific Computing" report (ViSC) that warned about the "firehose of data" that was resulting from computational experiments and electronic sensing [1]. The solution proposed by the ViSC report was to use visualisation to quickly spot patterns in the data that could then be used to guide the investigation towards hypotheses more likely to yield positive results. By using these "intuitive perceptual qualities as a basis for evaluation, verification and understanding", the ViSC panellists intended to put "the neurological machinery of the visual cortex to work".

In 2000, visualisation scientist Colin Ware concisely summed up the main advantages of modern visualization techniques [2]. As mentioned above, visualisation permits the apprehension of large amounts of data. The flexibility of human vision can perceive emergent properties such as subtle patterns and structures. It can compare small scale and large scale features at the same time. It can also help with the discernment of artefacts or mistakes in the gathering of the data itself. Yet despite these activities being at the intuitive level it is still possible to use them to suggest more formal hypotheses about the data in question. The early criticism that "pictures don't prove anything" has gradually been mitigated by the promise that apparent relationships can be later confirmed by applying more exact analytical methods.

Visualisations are created for people rather than for machines - they imply that not all computational processes can be fully automated and left to run themselves. Somewhere along the line a human being will need to evaluate or monitor the progress of the computation and make decisions about what to do next. Yet despite the fact that the material operations of software and data processing are perfectly objective and describable, they are rarely directly accessible to us. One of the fundamental properties of software is that once it is being executed it takes place on such a fine temporal and symbolic scale and across such a vast range of quantities of data that it has an intrinsically different materiality to that with which we are able to deal with unaided. Visualisation is one of the few techniques available for overcoming this distance. In the visualisation process, the transformations that lead from data to digital image are defined through software itself, often in a direct or 'live' relationship to it, yet aim to be apprehended at a level of human sensibility far beyond it. A visualisation is therefore distinguished by its algorithmic dependence on its source data and its perceptual independence from it.



Passamoquoddy Bay, oceanographic visualisation, 2000. Image courtesy of Colin Ware based on data provided by the Canadian Hydrographic Service.


Early writers on visualisation such as Edward Tufte developed guidelines and examples for how to design information graphics that are still influential today. Tufte's main concern is now referred to as the principle of being "expressive" - to remove all unnecessary graphical ornamentation and show as much data as possible - to "let the data speak for itself" [3]. To some extent, when we use computer graphics we can often "express" so much data that we do not have to choose which is the most significant. But even if we are able to show everything we may still do not know how to show it - how do we order the variables into an image in a way that expresses their interrelationships? The semiologist Jacques Bertin did important early work during the seventies on how to organise a "visual structure" that reflects the features and relations between the data [4]. The usual approach is to start from some basic knowledge about the data's internal structure. In theory the data we start with is raw and uninterpreted but in practice there is always some additional information about its composition, usually derived from the means by which it was gathered. For instance, if the data has up to three variables it can be directly mapped into a three dimensional graph of x, y, z values (or by transforming it using an intermediate stage called a 'data table'). Ware provides a typical example of such a visualisation from oceanography - a multibeam echo sounder scanning of the tides at Passamoquoddy Bay in Canada which produces a three dimensional array of height fields, rendered as a colour image [5]. This data used to be sampled and rendered as a set of contour lines, but the continuous computer generated image allows us to clearly see the more subtle features, textures and artefacts in all the millions of measurements made. Of course, we do not have to render it in this way - if we chose to we could unravel the array into a one dimensional sequence of values, interpret each one as a frequency and "play" the data as a series of tones. But this would be to ignore the variables' positional structure and we would almost certainly not be able to "see" the ripples and pockmark patterns that we can in the image. Ordering the values into a linear sequence might also imply precedence or ranking not in the original data. The internal structure of the data is spatial rather than temporal.

If we are using visualisation to forage for particular pieces of information such as which stocks are rising most steeply or in creating a graphic notation for structuring conceptual propositions, then we are dealing with more explicit functions of data catered for by specialised fields of information visualisation, “data mining” or knowledge visualisation. These disciplines are often closer to interface design, employing popular techniques such as interactive “fisheye” views, “table lens” document graphs or spatial “mind mapping” tools [6]. But in a more general context, if the properties of the data are yet to be discovered, then visualisation is less to do with retrieval, monitoring or communication and is more of an experimental technique. In contrast to a diagram that is constructed on the basis of a pre-established set of significances, a visualisation is about finding connections (or disconnections) between dataset attributes like amounts, classes or intervals which were previously unknown. Visualisations are always partial and provisional and they may entail the application of a number of different methods until the data gives up its secrets. The images frequently exhibit the continuous qualities of the familiar visual world despite the fact that they are utterly constructed. It is these implicit visual properties that are valued for their openness to perceptual inference, a continuous interplay of surface features rather than discrete graphic elements or symbols. At this end of the spectrum, visualisation is non-representational because it is speculatively mapped from the raw, isolated data or equations and not with respect to an already validated representational relation. A visualisation is instead a means to representation.

As recently as 2004, visualisation scientist Chaomei Chen described visualisation as still being an art rather than a science [7]. There is still no taxonomy of techniques that might help designers select one that is more effective for their requirements and no generic criteria with which to assess the value of a visualisation once they have. In the absence of guidance, there has been a tendency by some scientists to seize upon the underlying code of a successful visualisation and make it a de facto standard. Colin Ware has tried to remedy this by grounding visualisation as a specifically scientific discipline by combining the fields of physiology, human perception and cognitive studies [8]. This feeds into a desire amongst many scientists to conform visualisation to scientific method by treating visual perception and cognition in terms of computation itself, to be harnessed as an instrumental resource. For instance, the ability of the eye to instantly see that one visual feature is bigger than another is referred to as “computational offloading” in some places – “a diagram may allow users to avoid having to explicitly compute information because users can extract information ‘at a glance’” [9]. There is now a push to try to streamline visualisation by designing it for the faster “automatic processing” stage of human vision that deals with the unconscious detection of light, pattern, orientation and movement. If the abilities of this retinal level of processing can be defined and standardised then the hope is that visualisation can be freed of any inefficient dependence on learnt visual conventions, that it can promise a fast and universal “understanding without training” that crosses all cultural boundaries [10].

In all the literature there is little emphasis on how to see visualisations, only on how they are seen. Despite the fact that low level perceptual mechanics may not be formally “learnt”, they can still be exercised, sensitised, tuned and focused as an acquired skill. The editor of a film can see a hair on an individual frame that appears far too briefly for his audience to be conscious of. Visualisation as a practice is not just a question of designing for human perception but of ‘being perceptive’. In fact


some people's eyes have been "retrained" by visualisation itself until it has altered their apprehension of the world. One of the most early and ubiquitous forms of scientific visualisation were the images of fractals, chaos theory and complex systems of the late eighties [11]. Despite the fact that media theorists such as Vilem Flusser pointed out that these pictures were "images of equations" rather than "images of the world", they were frequently used to model the appearance of natural phenomena such as mountains, plants and marble textures [12]. Some scientists working with fractal modelling such as Michael Barnsley found that after a while they began to "see" the rivers, trees and clouds around them in terms of fractal mathematics [13], internalising concepts of self-similarity and strange attractors until they had become a way of thinking and perceiving itself, as though turning the whole world into a "natural" visualisation. Both algorithm and sensory vision are thus finally rejoined in the cortex, in an endless circularity of computation and perception.



'Black Spleenwort Fern'. Generated from fractal iterative functions, Michael Barnsley, 1988 [13].

Visualisation is usually separated out as a tool for knowledge formation rather than a visual form of knowledge itself. Although forms such as 'analogical representation' (which preserves some structural features of the object such as visual resemblance) or 'enactive knowledge' (which is bound to actions such as a certain skill) are recognized as valid forms of knowledge, scientists still mainly characterize their aims in terms of 'conceptual knowledge' - that which can be symbolically represented or discursively expressed [14]. This causes some uncertainty in the status of visuality – the literature frequently switches between statements like "using vision to think", "using visual computation to think" and "visual sense making" [15]. Michel Foucault described a similar situation in his study of the origins of modern systems of knowledge at the end of the Renaissance [16]. He pointed out how the principle of "resemblance", which had previously been so important, became relegated to a preliminary stage on the borders of knowledge during the Enlightenment. This was despite the fact that at the dawn of representational knowledge as now, no order could be established between two things until a visual similarity had occasioned their comparison. The use of memory and imagination in the discovery of a latent resemblance is what makes the creation of knowledge possible. Whether such visual relations will continue to be restricted to the rudimentary status of perceptual 'pre-processing' under the reign of visualisation will define one of the most important characteristics of knowledge in the age of computer software.


Although initially applied to imagery, visualisation has now become a more generic term that covers the sensory presentation of data and processing using interactive techniques, animation, sonics, haptics and multi-user VR environments. Over the course of the 1990s, visualisation has spread from the sciences into engineering disciplines, marketing, law, policy making and art and entertainment, indeed to any field that has found its object of interest replaced by datasets or computer models. It helps make visible the fluctuations in the international money market, defends the innocent through accident reconstruction, discloses network traffic to detect telephone fraud or reports the proportion of files consolidated by one's disk defragmenter.



Christian Nold, detail from "Greenwich Emotion Map", 2006 [17].

These new fields obviously exceed the original scientific aims of visualisation, yet even in art and design applications some form of cognitive knowledge may still be at issue. Christian Nold is an artist who has been building "bio maps" of communities using a mixture of consciously and unconsciously recorded data [17]. For the "Greenwich Emotion Map", 2006, groups of local residents each received a Galvanic Skin Response unit which measured their emotional arousal as they went for a walk around the neighbourhood. Every four seconds their level of excitation was recorded along with their geographical location as they reacted or failed to react to whatever coincidence of encounters, sights, and smells the city channelled to them that day. When they returned, their data was uploaded and plotted onto a map of Greenwich and annotated with written notes and photos they made at the time. When rendered as an overlay of nervous peaks and troughs, markers and pop-ups over a Google Earth satellite image, we are able to pick apart Google's naturalistic photographic image of Greenwich in terms of a mass of individual responses and rejoinders. The antecedent of this work might be the first photographic "visualisation" of a crowd taken in 1848 at the Chartist Demonstrations in London, which was used by a worried state to picture the political protest as a visibly quantifiable object of menace. In comparison, "Bio Mapping" recreates the urban crowd using data visualisation to become a visually dynamic object of fluctuating emotional intensities, informal commentaries and subjective trajectories.

There also exist many non-cognitive “visualisations” in common use. In some cases this is because they move so far from their source data that it disappears from relevance entirely. For example, it is easy to take any arbitrary data including random, unstructured data and contrive a rich pattern from it using elaborate visualisation tools. Noise functions are widely used in media production software as the starting point for synthetic image generation. By repeatedly applying a barrage of frequency filters, scalings and interpolation methods it is routinely possible to design the convincing appearance of natural phenomena such as marble, wood, smoke or fire, or the vertiginous synaesthetic abstractions familiar to users of the Windows Media Player. Due to this loss of significance of their source data, in these cases we move away from “data visualisation” as such to the more general category of computer generated “visualisation”.



Graham Harwood, “Lungs: Slave Labour”, 2005 [18].

But there are other non-cognitive visualisations whose power is derived from the very strain of stretching yet maintaining a connection to their original database. “Lungs: Slave Labour”, 2005 by Graham Harwood, is an acoustic visualisation based on a Nazi database of the foreign labourers that were forced to work in the ex-munitions factory that now houses the Centre for Media Art in Karlsruhe [18]. By interrogating their age, sex and height, “Lungs” is able to calculate their vital lung capacity and emit a ‘breath’ of air for each worker through a speaker system. The general aim of the “Lungs” project is to take computer records of local events or communities that have been reduced or demeaned to the status of information and to allow the people to re-experience and recover their own value. This attempt to give a database a pair of lungs reconnects people with a political atrocity in a very visceral way.

This last example might be seen as highly tendentious, but factually elaborates the politics involved in any representation of data. It still meets the central requirement for data visualisation of algorithmically deriving a sensory expression from the structures implicit in digital data, even when, and especially when, that expression takes us far from the realm of computer code. An effective visualisation makes use of

the greatest material distance between human senses and computer code achievable through the simplest material connections and procedures, to explore the most extreme perspectives which software can create of itself. It is its ability to put cognitive and affective modes of perception into creative tension with data structures as well as with each other, and to articulate the gap between the processing of data, social life and sensory experience that will allow visualisation to reach its full potential, both as a scientific and as an artistic technique.

1. Bill H. McCormick, Tom A. DeFanti, Maxine D Brown, "Visualisation in Scientific Computing". *Computer Graphics*, 21, no. 6, November 1987, U.S.A.
2. Colin Ware, *Information Visualisation: Perception for Design*, Morgan Kaufmann Publishers, San Francisco, 2000. p 2.
3. Edward Tufte, *Visual Explanations*, Graphics Press, U.S.A., 1997. p 45.
4. Jacques Bertin, *Graphics and Graphic Information Processing*, Walter de Gruyter, Berlin, New York, 1982 (1977).
5. Ware, p 2.
6. Sigmar-Olaf Tergan and Tanja Keller (eds). *Knowledge and Information Visualisation: Searching for Synergies (Lecture Notes in Computer Science)*, Springer-Verlag Berlin and Heidelberg GmbH & Co. K, 2005. p 5.
7. Chaomei Chen, *Information Visualisation*, Springer-Verlag UK; 2nd edition, 2004. p 1.
8. Ware, p 5.
9. M Scaife & Y Rogers, "External Cognition", *International Journal of Human-Computer Studies*, Vol 45, no. 2, pp 185 – 213. Academic Press, New York, 1996.
10. Ware, p 10.
11. Hans Otto Peitgen and Peter Richter, *The Beauty of Fractals: Images of Complex Dynamical Systems*, Springer-Verlag, Berlin, 1985.
12. Vilém Flusser, "Curies children: Vilém Flusser on an unspeakable future", *Artforum*, New York, March 1990.
13. Michael Barnsley, *Fractals Everywhere : The First Course in Deterministic Fractal Geometry*. Academic Press, New York, 1988. p 1.
14. Tergan and Keller, p 4.

15. Stuart Card, Jock Mackinlay, Ben Schneiderman (eds), Readings in Information Visualisation: Using Vision to Think. Morgan Kaufmann Publishers, San Francisco, 1999. pp 33, 34, 579.
16. Michel Foucault, The Order of Things, Routledge, London 2000 (1966). pp 67 – 68.
17. Christian Nold, “Greenwich Emotion Map”, 2006. Online at: <http://www.emotionmap.net>
18. Graham Harwood, “Lungs: Slave Labour”, 2005. Permanent collection, ZKM, Karlsruhe, Germany. <http://www.mongrel.org.uk/lungs>