



## What is visualisation?

Lev Manovich

To cite this article: Lev Manovich (2011) What is visualisation?, Visual Studies, 26:1, 36-49, DOI: [10.1080/1472586X.2011.548488](https://doi.org/10.1080/1472586X.2011.548488)

To link to this article: <https://doi.org/10.1080/1472586X.2011.548488>



Published online: 14 Mar 2011.



Submit your article to this journal [↗](#)



Article views: 2195



View related articles [↗](#)



Citing articles: 21 View citing articles [↗](#)

# What is visualisation?

LEV MANOVICH

*This article proposes that the practice of information visualisation (infovis), from its beginnings in the second part of the eighteenth century until today, relied on two key principles. The first principle is reduction. Infovis uses graphical primitives such as points, straight lines, curves and simple geometric shapes to stand in for objects and relations between them. The second principle is the use of spatial variables (position, size, shape and, more recently, movement) to represent key differences in the data and reveal patterns and relations. Following this analysis, the author discusses a more recent visualisation method which we can call 'direct visualisation' (or 'media visualisation'): creating new visual representations from the actual visual media objects (images, video) or their parts. The article analyses the well-known examples of artistic visualisations that use this method: Listening Post (Ben Rubin and Mark Hansen), Cinema Redux (Brendan Dawes) and Preservation of Selected Traces (Ben Fry). It further suggests that direct visualisation is particularly relevant for humanities, media studies and cultural institutions. Using the actual visual artefacts in visualisation as opposed to representing them by graphical primitives helps the researcher to understand meaning and/or cause behind the patterns she may observe, as well as to discover additional patterns. To illustrate this idea, examples of projects created in the author's lab at UCSD ([softwarestudies.com](http://softwarestudies.com)) are presented. Founded in 2007, the lab works on techniques and software to allow interactive exploration of large sets of visual cultural data using a direct visualisation approach and supervisualisation systems such as 215 megapixel HIPerSpace. The examples of its work are visualisations of 4553 covers of every issue of Time magazine published between 1923 and 2009; visualisations of all pages of every issue of Science and Popular Science magazines published between 1872 and 1922; and a set of visualisations of 1 million pages on Manga series.*

I first drew the Chart in order to clear up my own ideas on the subject, finding it very troublesome to retain a distinct notion of the changes that had taken place. I found it answered the purpose beyond my expectation, by bringing into one view the result of details that are dispersed over a very wide and intricate field of universal history; facts sometimes

connected with each other, sometimes not, and always requiring reflection each time they were referred to . . . [ In reference to] 'The Chart, No. 1', representing the rise and fall of all nations or countries that have been particularly distinguished for wealth or power, is the first of the sort that ever was engraved, and has, therefore, not yet met with public approbation.

(William Playfair, *An Inquiry into the Permanent Causes of the Decline and Fall of Powerful and Wealthy Nations* [1805])

The pretty photographs we and other tourists made in Las Vegas are not enough. How do you distort these to draw a meaning for a designer? How do you differentiate on a plan between form that is to be specifically built as shown and that which is, within constraints, allowed to happen? How do you represent the Strip as perceived by Mr. A. rather than as a piece of geometry? How do you show quality of light – or qualities of form – in a plan at 1 inch to 100 feet? How do you show fluxes and flows, or seasonal variation, or change with time?

(Venturi, Izenour, and Scott Brown 1972; emphasis in the original)

The 'Whole' is now nothing more than a provisional visualization which can be modified and reversed at will, by moving back to the individual components, and then looking for yet other tools to regroup the same elements into alternative assemblages.

(Latour 2009)

Information visualization is becoming more than a set of tools, technologies and techniques for large data sets. It is emerging as a medium in its own right, with a wide range of expressive potential.

(Eric Rodenbeck [Stamen Design], in keynote lecture at *Emerging Technology* conference, March 2008)

Visualization is ready to be a mass medium.

(Fernanda B. Viégas and Martin Wattenberg, in an interview for [infosthetics.com](http://infosthetics.com), May 2010)

---

Lev Manovich is the author of *Software Takes Command* (released under CC licence, 2008), *Soft Cinema: Navigating the Database* (MIT Press, 2005) and *The Language of New Media* (MIT Press, 2001), which is hailed as 'the most suggestive and broad-ranging media history since Marshall McLuhan'. Manovich is a professor in the Visual Arts Department, University of California-San Diego, a director of the Software Studies Initiative at California Institute for Telecommunications and Information Technology (CALIT2), and a professor at the European Graduate School.

In 2010 the Museum of Modern Art in New York presented a dynamic visualisation of its collection on five screens created by Imaginary Forces. The *New York Times* regularly features custom visualisations both in its printed and web editions created by the in-house NYTimes interactive team. The web is crawling with numerous sophisticated visualisation projects created by scientists, designers, artists and students. If you search for certain types of public data, the first result returned by Google search links to an automatically created interactive graph of these data. If you want to visualise your own data set, Many Eyes, Tableau Public and other sites offer free visualisation tools. It seems that 300 years after William Playfair's amazement at the cognitive power of information visualisation, others are finally acknowledging the projective power of his method.

What is information visualisation? Despite the growing popularity of infovis (a common abbreviation for 'information visualisation'), it is not so easy to come up with a definition which would work for all kinds of infovis projects being created today, and at the same time that would clearly separate it from other related fields such as scientific visualisation and information design. So let us start with a provisional definition that we can modify later. Let us define information visualisation as a mapping between discrete data and a visual representation. We can also use different concepts besides 'representation', each bringing an additional meaning. For example, if we believe that the brain uses a number of distinct representational and cognitive modalities, we can define infovis as a mapping from other cognitive modalities (such as mathematical and propositional) to an image modality.

My definition does not cover all aspects of information visualisation – such as the distinctions between static, dynamic (i.e. animated) and interactive visualisation – the latter, of course, being most important today. In fact, most definitions of infovis by computer science researchers equate it with the use of interactive computer-driven visual representations and interfaces. Here are examples of such definitions: 'Information visualisation (InfoVis) is the communication of abstract data through the use of interactive visual interfaces' (Keim et al. 2006); 'Information visualization utilizes computer graphics and interaction to assist humans in solving problems' (Purchase et al. 2008).

Interactive graphic interfaces in general, and interactive visualisation application in particular, bring all kinds of new techniques for manipulating data elements – from the ability to change how files are shown on the desktop in modern operating systems to multiple coordinated views available in some visualisation software, such as

Mondrian.<sup>1</sup> However, regardless of whether you are looking at a visualisation printed on paper or a dynamic arrangement of graphic elements on your computer screen, which you generated using interactive software and can change at any moment, in both cases the image you are working with is a result of mapping. So what is special about images such mapping produces? This is the focus of my article.

For some researchers, information visualisation is distinct from scientific visualisation in that the latter uses numerical data while the former uses non-numeric data such as text and networks of relations.<sup>2</sup> Personally, I am not sure that this distinction holds in practice. Certainly, plenty of infovis projects use numbers as their primary data, but even when they focus on other data types, they still often use some numerical data as well. For instance, typical network visualisation may use both the data about the structure of the network (which nodes are connected to each other) and the quantitative data about the strength of these connections (for example, how many messages are exchanged between members of a social network). As a concrete example of infovis which combines non-numerical and numerical data, consider a well-known project, *History Flow* (Fernanda B. Viégas and Martin Wattenberg, 2003), which shows how a given Wikipedia page grows over time as different authors contribute.<sup>3</sup> The contribution of each author is represented by a line. The width of the line changes over time, reflecting the amount of text contributed by an author to the Wikipedia page. To take another infovis classic, *Flight Patterns* (Aaron Koblin, 2005) uses the numerical data about the flight schedules and trajectories of all planes that fly over the United States to create an animated map which displays the pattern formed by their movement over a 24-hour period.<sup>4</sup>

Rather than trying to separate information visualisation and scientific visualisation using some a priori idea, let us instead enter each phrase in Google image search and compare the results. The majority of images returned by searching for 'information visualisation' are two-dimensional and use vector graphics (i.e. points, lines, curves and other simple geometric shapes). The majority of images returned by searching for 'scientific visualisation' are three-dimensional; they use solid 3D shapes or volumes made from 3D points. The results returned by these searches suggest that the two fields indeed differ – not because they necessarily use different types of data, but because they privilege different visual techniques and technologies.

Scientific visualisation and information visualisation come from different cultures (science and design, respectively), and their development corresponds to

different areas of computer graphics technology. Scientific visualisation developed in the 1980s along with the field of 3D computer graphics, which at that time required specialised graphics workstations. Information visualisation developed in the 1990s along with the rise of desktop 2D graphics software and the adoption of PCs by designers; its popularity accelerated in the 2000s – the two key factors being the easy availability of big data sets via Application Programming Interfaces (APIs) provided by major social network services since 2005, and new high-level programming languages specifically designed for graphics (e.g. *Processing*)<sup>5</sup> and software libraries for visualisation (e.g. *Prefuse*).<sup>6</sup>

Can we differentiate information visualisation from information design? This is trickier, but here is my way of doing it. Information design starts with the data that already have a clear structure, and its goal is to express this structure visually. For example, the famous London tube map designed in 1931 by Harry Beck uses structured data: tube lines, tube stations and their locations over London geography.<sup>7</sup> In contrast, the goal of information visualisation is to discover the structure of a (typically large) data set. This structure is not known a priori; a visualisation is successful if it reveals this structure. A different way to express this is to say that information design works with information, while information visualisation works with data. As is always the case with the actual cultural practice, it is easy to find examples that do not fit such distinction – but a majority do. Therefore, I think that this distinction can be useful in allowing us to understand the practices of information visualisation and information design as partially overlapping but ultimately different in terms of their functions.

Finally, what about the earlier practices of visual display of quantitative information in the nineteenth and twentieth centuries that are known to many via the examples collected in the pioneering books by Edward Tufte (Tufte 1983, 1990, 1997, 2006)? Do they constitute infovis as we understand it today? As I already noted, most definitions provided by the researchers working within computer science equate information visualisation with the use of interactive computer graphics.<sup>8</sup> Using software, we can visualise much larger data sets than was possible previously; create animated visualisation; show how processes unfold in time; and, most importantly, manipulate visualisations interactively. These differences are very important – but for the purposes of this article, which is concerned with the visual language of infovis, they do not matter. When we switched from pencils to computers, this did not affect the core idea of visualisation – mapping some properties of the data into a visual representation.

Similarly, while availability of computers led to the development of new visualisation techniques (scatter plot matrix, treemaps, etc.), the basic visual language of infovis remained the same as it was in the nineteenth century – points, lines, rectangles and other graphic primitives. Given this continuity, I will use the term ‘infovis’ to refer to both earlier visual representations of data created manually and contemporary software-driven visualisation.

## REDUCTION AND SPACE

In my view, the practice of information visualisation from its beginnings in the second part of the eighteenth century until today relied on two key principles. The first principle is reduction. Infovis uses graphical primitives such as points, straight lines, curves and simple geometric shapes to stand in for objects and relations between them – regardless of whether these are people, their social relations, stock prices, income of nations, unemployment statistics, or anything else. By employing graphical primitives (or, to use the language of contemporary digital media, vector graphics), infovis is able to reveal patterns and structures in the data objects that these primitives represent. However, the price being paid for this power is extreme schematisation. We throw away 99 per cent of what is specific about each object to represent only 1 per cent – in the hope of revealing patterns across this 1 per cent of objects’ characteristics.

Information visualisation is not unique in relying on such extreme reduction of the world in order to gain new power over what is extracted from it. It comes into its own in the first part of the nineteenth century when, in the course of just a few decades, almost all graph types commonly found today in statistical and charting programs were invented.<sup>9</sup> This development of the new techniques for visual reduction parallels the reductionist trajectory of modern science in the nineteenth century. Physics, chemistry, biology, linguistics, psychology and sociology propose that both the natural and the social world should be understood in terms of simple elements (molecules, atoms, phonemes, just-noticeable sensory differences, etc.) and the rules of their interaction. This reductionism becomes the default ‘meta-paradigm’ of modern science and it continues to rule scientific research today. For instance, currently popular paradigms of complexity and artificial life focus our attention on how complex structures and behaviour emerge out of the interaction of simple elements.

Even more direct is the link between nineteenth-century infovis and the rise of social statistics. Philip Ball summarises the beginnings of statistics in this way:

In 1749 the German scholar Gottfried Achenwall suggested that since this ‘science’ [the study of society by counting] dealt with the natural ‘states’ of society, it should be called Statistik. John Sinclair, a Scottish Presbyterian minister, liked the term well enough to introduce it into the English language in his epic Statistical Account of Scotland, the first of the 21 volumes of which appeared in 1791. The purveyors of this discipline were not mathematicians, however, nor barely ‘scientists’ either; they were tabulators of numbers, and they called themselves ‘statists’. (Ball 2004, 64–5)

In the first part of the nineteenth century many scholars, including Adolphe Quetelet, Florence Nightingale, Thomas Buckle and Francis Galton, used statistics to look for ‘laws of society’. This inevitably involved summarisation and reduction – calculating the totals and averages of the collected numbers about citizens’ demographic characteristics, comparing the averages for different geographical regions, asking if they followed a bell-shaped normal distribution, etc. It is therefore not surprising that many – if not most – graphical methods that are standard today were invented during this time for the purposes of representations of such summarised data. According to Michael Friendly and Daniel J. Denis, between 1800 and 1850,

[i]n statistical graphics, all of the modern forms of data display were invented: bar and pie charts, histograms, line graphs and time-series plots, contour plots, and so forth. (Friendly and Denis 2010)

Do all these different visualisation techniques have something in common besides reduction? They all use spatial variables (position, size, shape and, more recently, curvature of lines and movement) to represent key differences in the data and reveal most important patterns and relations. This is the second (after reduction) core principle of infovis practice as it was practised for 300 years – from the very first line graphs (1711), bar charts (1786) and pie charts (1801) to their ubiquity today in all graphing software, such as Excel, Numbers, Google Docs and OpenOffice.<sup>10</sup>

This principle can be rephrased as follows: infovis privileges spatial dimensions over other visual dimensions. In other words, we map the properties of our data that we are most interested in onto a topology and geometry. Other less important properties of the objects are represented through different visual dimensions – tones, shading patterns, colours or transparency of the graphical elements.

As examples, consider two common graph types: a bar chart and a line graph. Both first appeared in William Playfair’s *Commercial and Political Atlas* (1786) and became commonplace in the early nineteenth century. A bar chart represents the differences between data objects via rectangles that have the same width but different heights. A line graph represents changes in the data values over time via changing height of the line.

Another common graph type – scatter plot – similarly uses spatial variables (positions and distances between points) to make sense of the data. If some points form a cluster, this implies that the corresponding data objects have something in common; if you observe two distinct clusters, this implies that the objects fall into two different classes; and so on.

Consider another example – network visualisations which function today as distinct symbols of ‘network society’ (see Manuel Lima’s authoritative gallery, [visualcomplexity.com](http://visualcomplexity.com), which currently houses over 700 network visualisation projects). Like bar charts and line graphs, network visualisations also privilege spatial dimensions: position, size and shape. Their key addition is the use of straight or curved lines to show connections between data objects. For example, in *distellamap*, Ben Fry (2005) connects pieces of code and data by lines to show the dynamics of the software execution in Atari 2600 games.<sup>11</sup> In Marcos Weskamp’s *Flickrgraph* (2005), the lines visualise the social relationships between users of flickr.com.<sup>12</sup> (Of course, many other visual techniques can also be used in addition to lines to show relations – see, for instance, a number of maps of science created by Katy Borner and her colleagues at Information Visualization Lab at Indiana University.)<sup>13</sup>

I believe that the majority of information visualisation practices from the second part of the eighteenth century until today follow the same principle – reserving spatial arrangement (we can call it ‘layout’) for the most important dimensions of the data, and using other visual variables for remaining dimensions. This principle can be found in visualisations ranging from the famous dense graphic showing Napoleon’s march on Moscow by Charles Joseph Minard (1869)<sup>14</sup> to the recent *The Evolution of The Origin of Species* by Stefanie Posavec and Greg McInerney (2009).<sup>15</sup> Distances between elements and their positions, shape, size, lines, curvature and other spatial variables code quantitative differences between objects and/or their relations (for instance, who is connected to whom in a social network).

When visualisations use colours, fill-in patterns or different saturation levels, typically this is done to partition graphic elements into groups. In other words,



these non-spatial variables function as group labels. For example, Google Trends uses line graphs to compare search volumes for different words or phrases; each line is rendered in a different colour.<sup>16</sup> However, the same visualisation could have simply used labels attached to the lines – without different colours. In this case, colour adds readability but it does not add new information to the visualisation.

The privileging of spatial over other visual dimensions was also true of plastic arts in Europe between the sixteenth and nineteenth centuries. A painter first worked out the composition for a new work in many sketches; next, the composition was transferred to a canvas and shading was fully developed in monochrome; only after that was colour added. This practice assumed that the meaning and emotional impact of an image depends most of all on the spatial arrangements of its parts, as opposed to colours, textures and other visual parameters. In classical Asian ‘ink and wash painting’ which first appeared in the seventh century in China and was later introduced to Korea and then Japan (in the fourteenth century), colour did not even appear. The painters used exclusively black ink exploring the contrasts between objects’ contours, their spatial arrangements and different types of brushstrokes.

It is possible to find information visualisations where the main dimension is colour – for instance, a common traffic light which ‘visualises’ the three possible behaviours of a car driver: stop, get ready, go. This example shows that if we fix spatial parameters of visualisation, colour can become the salient dimension. In other words, it is crucial that the three lights have exactly the same shape and size. Apparently, if all elements of the visualisation have the same values on spatial dimensions, our visual system can focus on the differences represented by colours, or other non-spatial variables.

Why do visualisation designers – be they the inventors of graph and chart techniques at the end of the eighteenth and early nineteenth centuries, or millions of people who now use these graph types in their reports and presentations, or the authors of more experimental visualisations featured on [infoaesthetics.com](http://infoaesthetics.com) and [visualcomplexity.com](http://visualcomplexity.com) – privilege spatial variables over other kinds of visual mappings? In other words, why are colour, tone, transparency and symbols used to represent secondary aspects of data while the spatial variables are reserved for the most important dimensions? Without going into the rich but still very incomplete knowledge about vision accumulated by neuroscience and experimental psychology, we can make a simple guess. The creators of visualisations follow human visual

perception that also privileges spatial arrangements of parts of a scene over its other visual properties in making sense of this scene. Why would the geometric arrangement of elements in a scene be more important to human perception than other visual dimensions? Perhaps this has to do with the fact that each object occupies a unique part of the space. Therefore it is crucial for a brain to be able to segment a 3D world into spatially distinct objects which are likely to have distinct identities (people, sky, ground, cards, buildings, etc.). Different object types can also be often identified with unique 2D forms and arrangements of these forms. A tree has a trunk and branches; a human being has a head, a torso, arms and legs; etc. Therefore identifying 2D forms and their arrangements is also likely to play an important role in object recognition.

An artist or a designer may pay more attention to other visual properties of a scene, such as textures and rhythms of colour (think of twentieth-century art) – but in an everyday perception, spatial properties are what matters most. How close two people are to each other; the expression on their faces; their relative size which allows the observer to estimate his distance from her; the characteristic shapes of different objects which allows her to recognise them – all these and many other spatial characteristics which our brains instantly compute from the retinal input are crucial for our daily existence.

I think that this key of spatial variables for human perception may be the reason why all standard techniques for making graphs and charts that developed from the eighteenth to the twentieth century use spatial dimensions to represent the key aspects of the data, and reserve other visual dimensions for less important aspects. However, we should also keep in mind the evolution of visual display technologies, which constrain what is possible at any given time. Only in the 1990s, when people started using computers to design and present visualisations on computer monitors, did colour become the norm. Colour printing is still significantly more expensive than using a single colour – so even today science journals are printed in black and white. Thus, the extra cost associated with creating and printing colour graphics during the last two centuries was probably an important factor responsible for the privileging of spatial variables.

When colour, shading and other non-spatial visual parameters were used in visualisations created in the nineteenth century and most of the twentieth century, they usually represented only a small number of discrete values – that is, they acted as ‘categorical variables’. However, today the fields of computer-based scientific

visualisation, geovisualisation and medical imaging often use such parameters with much larger scales. Since today computers commonly allocate 8 bits to store values for each of red, green and blue channels, computer monitors can show 16 million unique colours. Therefore colour, shading and transparency are now commonly employed in these fields to show continuously varying qualities such as temperature, gas density, elevation and gravity waves. Does this not contradict my statement that spatial arrangement is key to information visualisation?

We can solve this puzzle if we take into account a fundamental difference between information visualisation and scientific visualisation/geovisualisation which I have not yet mentioned. Infovis uses arbitrary spatial arrangements of elements to represent the relationships between data objects. Scientific and medical visualisation and geovisualisation typically work with an *a priori* fixed spatial layout of a real physical object such as a brain, a coastline, a galaxy, etc. Since the layout in such visualisations is already fixed and cannot be arbitrarily manipulated, colour and/or other non-spatial parameters are used instead to show new information. A typical example of this strategy is a *heat map* which uses colour hue and saturation to overlay information over a spatial map.<sup>17</sup>

The two key principles that I suggested – *data reduction* and *privileging of spatial variables* – do not account for all possible visualisations produced during last 300 years. However, they are sufficient to separate infovis (at least as it was commonly practised until now) from other techniques and technologies for visual representation: maps, engraving, drawing, oil painting, photography, film, video, radar, MRI, infrared spectroscopy, etc. They give infovis its unique identity – the identity which remained remarkably consistent for almost 300 years (i.e. until the 1990s).

## VISUALISATION WITHOUT REDUCTION

The meanings of the word ‘visualise’ include ‘make visible’ and ‘make a mental image’. This implies that until we ‘visualise’ something, this ‘something’ does not have a visual form. It becomes an image through a process of visualisation.

If we survey the practice of infovis from the eighteenth century until the end of the twentieth century, the idea that visualisation takes data that is not visual and maps it into a visual domain indeed works quite well. However, it seems to no longer adequately describe certain new visualisation techniques and projects developed since the middle of the 1990s. Although these techniques and projects are commonly discussed as ‘information

visualisation’, is it possible that they actually represent something else – a fundamentally new development in the history of representational and epistemological technologies, or at least a new broad visualisation method for which we don’t yet have an adequate name?

Consider a technique called tag cloud.<sup>18</sup> The technique was popularised by Flickr in 2005 and today it can be found on numerous web sites and blogs. A tag cloud shows most common words in a text in the font size corresponding to their frequency in the text.

We can use a bar chart with text labels to represent the same information – which in fact may work better if the word frequencies are very similar. But if the frequencies fall within a larger range, we don’t have to map the data into a new visual representation such as the bars. Instead, we can vary the size of the words themselves to represent their frequencies in the text.

Tag cloud exemplifies a broad method that can be called media visualisation: creating new visual representations from the actual visual media objects, or their parts. Rather than representing text, images, video or other media through new visual signs such as points or rectangles, media visualisations build new representations out of the original media. Images remain images; text remains text.

In view of our discussion of the data reduction principle, we can also call this method direct visualisation, or visualisation without reduction. In direct visualisation, the data are reorganised into a new visual representation that preserves its original form. Usually, this does involve some data transformation, such as changing data size. For instance, tag cloud reduces the size of text to a small number of most frequently used words. However, this is a reduction that is quantitative rather than qualitative. We do not substitute media objects with new objects (i.e. graphical primitives typically used in infovis) which only communicate selected properties of these objects (for instance, bars of different lengths representing word frequencies). My phrase ‘visualisation without reduction’ refers to this preservation of a much richer set of properties of data objects when we create visualisations directly from them.

Not all direct visualisation techniques, such as tag cloud, originated in the twenty-first century. If we project this concept retroactively into history, we can find earlier techniques that use the same idea. For instance, a familiar book index can be understood as a direct visualisation technique. Looking at a book’s index, one can quickly see if particular concepts or names are

important in the book – they will have more entries; less important concepts will take up only a single line.

While both book index and tag cloud exemplify direct visualisation, it is important to consider the differences between them. The older book index technique relied on the typesetting technology used for printing books. Since each typeface was only available in a limited number of sizes, the idea that you can precisely map the frequency of a particular word onto its font size was counter-intuitive – so it was not invented. In contrast, the tag cloud technique is a typical expression of what we can call ‘software thinking’ – that is, the ideas that explore the fundamental capacities of modern software. Tag cloud explores the capacities of software to vary every parameter of a representation and to control it using external data. The data can come from a scientific experiment, from a mathematical simulation, from the body of the person in an interactive installation, from calculating some properties of the data, and so on. If we take these two capacities for granted, the idea of arbitrarily changing the size of words based on some information – such as their frequency in a text – is something we may expect to be ‘actualised’ in the process of cultural evolution. (In fact, all contemporary interactive visualisation techniques rely on the same two fundamental capacities.)

The rapid growth in the number and variety of visualisation projects, software applications and web services since the late 1990s was enabled by the advances in computer graphics capacities of personal computers, including both hardware (processors, RAM, displays) and software (C and Java graphics libraries, Flash, Processing, Flex, Prefuse, etc.). These developments both popularised information visualisation and also fundamentally changed its identity by foregrounding animation, interactivity and also more complex visualisations, which represent connections between many more objects.<sup>19</sup> But along with these three highly visible trends, the same advances also made possible the ‘direct visualisation’ approach – although it has not been given its own name so far.

## DIRECT VISUALISATION: EXAMPLES

Let us discuss three well-known projects which exemplify ‘direct visualisation’: *Listening Post*, *Cinema Redux* and *Preservation of Selected Traces*.<sup>20</sup>

*Cinema Redux* was created by interactive designer Brendan Dawes in 2004.<sup>21</sup> Dawes wrote a program in processing that sampled a film at the rate of one frame per second and scaled each frame to 8 x 6 pixels. The program then arranged these minute frames in a

rectangular grid with every row representing a single minute of the film. Although Dawes could have easily continued this process of sampling and remapping – for instance, representing each frame through its dominant colour – he chose instead to use the actual scaled-down frames from the film. The resulting visualisation represents a trade-off between the two possible extremes: preserving all the details of the original artefact, and abstracting its structure completely. Higher degrees of abstraction may make the patterns in cinematography and narrative more visible, but would also remove the viewer further from the experience of the film. Staying closer to the original artefact preserves the original detail and aesthetic experience, but may not be able to reveal some of the patterns.

What is important in the context of our discussion is not the particular parameters which Dawes used for *Cinema Redux*, but that he reinterpreted the previous constant of visualisation practice as a variable. Previously infovis designers mapped data into new diagrammatic representation consisting of graphical primitives. This was the default practice. With computers, a designer can select any value on the ‘original data’/abstract representation dimension. In other words, a designer can now choose to use graphical primitives, or the original images exactly as they are, or any format in between. Thus, while the project’s title refers to the idea of reduction, in the historical content of earlier infovis practice it can be actually understood as expansion – that is, expanding typical graphical primitives (points, rectangles, etc.) into the actual data objects (film frames).

Before software, visualisation usually involved the two-stage process of first counting or quantifying data, and then representing the results graphically. Software allows for direct manipulation of the media artefacts without quantifying them. As demonstrated by *Cinema Redux*, these manipulations can make visible the relations between a numbers of artefacts. Of course, such visualisation without quantification is made possible by the a priori quantification required to turn any analogue data into a digital representation. In other words, it is the ‘reduction’ first performed by the digitisation process which paradoxically now allows us to visualise the patterns across sets of analogue artefacts without reducing them to graphical signs.

For another example of direct visualisation, let us turn to Ben Fry’s *Preservation of Selected Traces*.<sup>22</sup> This web project is an interactive animation of the complete text of Darwin’s *Evolution of the Species*. Fry uses different colours to show the changes made by Darwin in each of six editions of his famous book. As the animation plays, we see the evolution of the book’s text from edition to





FIGURE 1. Visualisations of the text and patterns in Tolstoy's *Anna Karenina*. Data: Project Gutenberg. If a line contains the word 'Anna' it appears in blue. If a line contains the word 'Vronsky' it appears in red. In this version, each line of text is rendered over rectangles which extend the length of the line. Created by Software Studies Initiative, Calit2/UCSD (softwarestudies.com).

edition, with sentences and passages deleted, inserted and re-written. In contrast to typical animated information visualisations, which show some spatial structure constantly changing its shape and size in time, reflecting changes in the data (for example, changing structure of a social network over time), in Fry's project the rectangular shape containing the complete text of Darwin's book always stays the same – what changes is its content. This allows us to see how over time the pattern of additions and revisions become more and more intricate, as the changes from all the editions accumulate.

At any moment in the animation we have access to the complete text of Darwin's book – as opposed to only diagrammatic representation of the changes. At the same time, it can be argued that that *Preservation of Selected Traces* does involve some data reduction. Given the typical resolution of computer monitors and web bandwidth today, Fry was not able to actually show all the actual book text at the same time.<sup>23</sup> Instead sentences are rendered as tiny rectangles in different colours. However, when you mouse over any part of the image, a pop-up window shows the actual text. Because all the text of Darwin's book is easily accessible to the user in this way, I think that this project can be considered an example of direct visualisation.

Finally let us add one more example – *Listening Post* by Ben Rubin and Mark Hansen (2001).<sup>24</sup> Usually this work is considered to be a computer-driven installation – rather than an example of infovis. *Listening Post* pulls text fragments from online chat rooms in real time, based on various parameters set by the authors, and streams them across a display wall made from a few hundred small screens in a six-act looping sequence. Each act uses its own distinct spatial layout to arrange dynamically changing text fragments. For instance, in one act, the phrases move across the wall in a wave-like pattern; in another act, words appear and disappear in a checkerboard pattern. Each act also has its own distinct sound environment driven by the parameters extracted from the same text that is being animated on the display wall.

One can argue that *Listening Post* is not a visualisation because the spatial patterns are pre-arranged by the authors and not driven by the data. This argument makes sense – but I think it is important to keep in mind that while layouts are pre-arranged, the data in these layouts are not – they are a result of the real-time data mining of the web. So while the text fragments are displayed in pre-defined layouts (wave, checkerboard, etc.), because the content of these fragments is always different, the overall result is also always unique. Note

that if the authors were to represent the text via abstract graphical elements, we would simply end up with the same abstract pattern in every repetition of an act. But because they show the actual text that changes all the time, the patterns that emerge inside the same layout are always different.

This is why I consider *Listening Post* to be a perfect representation of direct visualisation – the patterns it presents depend as much on what all the text fragments which appear on the screen wall actually say as on their pre-defined composition. We can find other examples of info projects that similarly flow the data into pre-defined layouts. Manuel Lima identified what he calls a ‘syntax’ of network visualisations – commonly used layouts such as radial convergence, arc diagrams, radial centralised networks and others.<sup>25</sup> The key difference between most of these network visualisations and *Listening Post* lies in the fact that the former often rely on the existing visualisation layout algorithms. Thus they implicitly accept ideologies behind these layouts – in particular the tendency to represent a network as a highly symmetrical and/or circular structure. The authors of *Listening Post* wrote their own layout algorithms that allowed them to control the layouts’ intended meanings. It is also important that they use six very different layouts that cycle over time. The meaning and aesthetic experience of this work – showing both the infinite diversity of the web and at the same time the existence of many repeating patterns – to a significant extent derive from the temporal contrasts between these layouts. Eight years before Bruno Latour’s article (quoted at the beginning) – where Latour argues that our ability to create ‘a provisional visualisation which can be modified and reversed’ allows us to think differently since any ‘whole’ we can construct now is just one of numerous others – *Listening Post* beautifully staged this new epistemological paradigm enabled by interactive visualisation.

The three influential projects considered above demonstrate that in order to highlight patterns in the data we don’t have to reduce them by representing data objects via abstract graphical elements. We also do not have to summarise the data, as is common in statistics and statistical graphics – think, for instance, of a histogram which divides data into a number of bins. This does not mean that in order to qualify as a ‘direct visualisation’ an image has to show all 100% of the original data – every word in a text, every frame in a movie. Out of the three examples I just discussed, only *Preservation of Selected Traces* does this. Both *Cinema Redux* and *Listening Post* do not use all the available data – instead, they sample them. The first project samples a feature film at the fixed rate of one frame per second; the

second project filters the online conversations using set criteria that change from act to act. However, what is crucial is that the elements of these visualisations are not the result of the remapping of the data into some new representation format – they are the original data objects selected from the complete data set. This strategy is related to the traditional rhetorical figure of synecdoche – specifically, its particular case where a specific class of thing refers to a larger, more general class.<sup>26</sup> (For example, in *Cinema Redux* one frame stands for a second of a film.)

While sampling is a powerful technique for revealing patterns in the data, *Preservation of Selected Traces* demonstrates that it is also possible to reveal patterns while keeping 100 per cent of the data. But you have already been employing this strategy if you have ever used a magic marker to highlight important passages of a printed text. Although text highlighting is not normally thought of, we can see that in fact it is an example of ‘direct visualisation without sampling’.

*Cinema Redux* and *Preservation of Selected Traces* also break away from the second key principle of traditional visualisation – communication of meaning via spatial arrangements of the elements. In both projects, the layout of elements is dictated by the original order of the data – shots in a film, sentences in a book. This is possible and also appropriate because the data they visualise are not the same as the typical data used in infovis. A film or a book is not just a collection of data objects – they are narratives made from these objects (i.e. the data have a sequential order). Although it is certainly possible to create effective visualisations that remap a narrative sequence into a completely new spatial structure, as in *Listening Post* (see, for instance, *Writing Without Words* by Stefanie Posavec [2008]<sup>27</sup> and *The Shape of Song* by Martin Wattenberg [2001]),<sup>28</sup> *Cinema Redux* and *Preservation of Selected Traces* demonstrate that preserving the original sequences is also effective.

Preserving the original order of data is particularly appropriate in the case of cultural data sets that have a time dimension. We can call such data sets ‘cultural time series’. Whether it is a feature film (*Cinema Redux*), a book (*Preservation of Selected Traces*) or a long Wikipedia article (*History Flow*), the relationships between the individual elements (film shots, book’s sentences) and also between larger parts of a work (film scenes, book’s paragraphs and chapters) separated in time are of primary importance to the work’s evolution and meaning and its experience by the users. We consciously or unconsciously notice many of these patterns during watching/reading/interacting with the work and



projecting time into space – laying out movie frames, book sentences, magazine pages in a single image – gives us new possibilities to study them. Thus, *space* turns out to play a crucial role in direct visualisation after all: it *allows us to see patterns between media elements that are normally separated by time*.

Let me add to this discussion a few more examples of direct visualisation that my students and I created at my lab – Software Studies Initiative (softwarestudies.com).<sup>29</sup> Inspired by the artistic projects which pioneered the direct visualisation approach, as well by the resolution and real-time capabilities of supervisualisation interactive systems such as HIPerSpace (35,840 x 8,000 pixels, 286,720,000 pixels in total)<sup>30</sup> developed at California Institute for Telecommunication and Information (Calit2)<sup>31</sup> where our lab is located, my group has been working on techniques and software to allow interactive exploration of large sets of visual cultural data. Some of the visualisations we created use the same strategy as *Cinema Redux* – arranging a large set of images in a rectangular grid. However, having access to a very high resolution display sometimes allows us to include all 100% of data – as opposed to having to sample them. For example, we created an image showing 4553 covers of every issue of *Time* magazine published between 1923 and 2009 (*Mapping Time*, Jeremy Douglass and Lev Manovich, 2009).<sup>32</sup> We also compared the use of images in *Science* and *Popular Science* magazines by visualising approximately 10,000 pages from each magazine during first decades of their publication (*The Shape of Science*, William Huber, Lev Manovich, Tara Zapel, 2010).<sup>33</sup> Our most data-intensive direct visualisation is 44,000 x 44,000 pixels; it shows 1,074,790 Manga pages organised by their stylistic properties (*Manga Style Space*, Lev Manovich and Jeremy Douglass, 2010).<sup>34</sup>

Like *Cinema Redux*, *Mapping Time* and *The Shape of Science* make equal the values of spatial variables to reveal the patterns in the content, colours and compositions of the images. All images are displayed at the same size arranged into a rectangular grid according to their original sequence. Essentially, these direct visualisation use only one dimension – with the sequence of images wrapped around into a number of rows to make it easier to see the patterns without having to visually scan very long images. However, we can turn such one-dimensional image timelines into 2D, with the second dimension communicating additional information. Consider a 2D timeline of *Time* covers we created (*Timeline*, Jeremy Douglass and Lev Manovich, 2009).<sup>35</sup> The horizontal axis is used to position images in the original sequence: time runs from left to right, and every cover is arranged according to its publication date.

The positions on the vertical axis represent new information – in this case, average saturation (the perceived intensity of colours) of every cover, which we measured using image analysis software.

Such mapping is particularly useful for showing variation in the data over time. We can see how colour saturation gradually increases during *Time*'s publication, reaching its peak in 1968. The range of all values (i.e. variance) per year of publication also gradually increases – but it reaches its maximum value a few years earlier. It is perhaps not surprising to see that the intensity (or 'aggressiveness') of mass media as exemplified by *Time* covers gradually increases up to the end of the 1960s, as manifested by changes in saturation and contrast. What is unexpected, however, is that since the beginning of the twenty-first century, this trend is reversed: the covers now have less contrast and less saturation.

The strategy used in this visualisation is based on the familiar technique – a scatter graph. However, if a normal scatter graph reduces the data displaying each object as a point, we display the data in their original form. The result is new graph type, which is literally made from images – that's why it is appropriate to call it an 'image graph'.<sup>36</sup>

## WHAT IS VISUALISATION?

In an article on the emerging practice of artistic visualisation written in 2002 I defined visualisation as 'a transformation of quantified data which is not visual is into a visual representation'. At that time I wanted to stress that visualisation participates in the reduction projects of modern science and modern art which led to the choice of the article's title: 'Data Visualisation as New Abstraction and Anti-Sublime'.<sup>37</sup> I think that this emphasis was appropriate given the types of infovis typically created at that time. (Although I used a somewhat different formulation for the definition that appears in the beginning of the present article – 'a remapping from other codes to a visual code' – the two definitions express the same idea.)

Most information visualisation today continues to employ graphical primitives. However, as the examples we looked at demonstrate, alongside this 'mainstream' infovis, we can find another trend – projects where the data being visualised are already visual – text, film frames, magazine covers. In other words, these projects create new visual representations of the original visual data without translating them into graphic signs. They also often break away from the second key principle of





FIGURE 2. *Mapping Time*. Jeremy Douglass and Lev Manovich, 2009. Data: The covers of every issue of *Time* magazine published from 1923 to summer 2009. Total number of covers: 4535. A large percentage of the covers included red borders. We cropped these borders and scaled all images to the same size to allow a user to see more clearly the temporal patterns across all covers.



infovis – mapping the most important dimensions in the data into spatial variables.

So does ‘direct visualisation’ constitute a form of infovis, or is it a different method altogether? We have two choices. Either we need to accept that this is something fundamentally different, or alternatively, we can revise our understanding of infovis.

Given that all direct visualisations we looked at aim to make visible patterns and relations in the data, this aim certainly aligns them with infovis as it developed during the last 300 years. It is also relevant to note that some of the most well-known infovis projects of the last 15 years follow a direct visualisation approach. This is true of *Cinema Redux* and *Preservation of Selected Traces* and other seminal projects which I have not discussed in detail, such as *Talmud Project* (David Small, 1999), *Valence* (Ben Fry, 2001) and *TextArc* (W. Bradford Paley, 2002). This means that people intuitively identify them as infovis even though they consist not of vector elements but of media (text or images). In another example, a *Phrase Net* technique which was developed by Frank van Ham, Martin Wattenberg and Fernanda Viégas and awarded ‘Best Paper’ at the IEEE InfoVis 2009 conference also operates within a direct visualisation paradigm (van Ham, Wattenberg, and Viégas 2009).

Does this mean that what we took to be the core principle of information visualisation during its first three centuries – reduction to graphic primitives – was only a particular historical choice, an artefact of the available graphics technologies? I think so. Similarly, the privileging of spatial variables over other visual parameters may also turn out to be a historically specific strategy – rather than the essential principle of infovis. The relatively new abilities brought by computer graphics to precisely control – that is, assign values within a large range – colour, transparency, texture and any other visual parameter of any part of an image allow us to start using these non-spatial parameters to represent the key dimensions of the data. This is already common in scientific and medical visualisation and geovisualisation – but not yet in information visualisation.

Why has infovis continued to rely on computer-generated vector graphics during the 1990s and 2000s when the speed with which computers can render images has been progressively increasing? Perhaps the main factor has been the focus on the World Wide Web as the preferred platform for delivering interactive visualisation. The web technologies made it relatively easy to create vector graphics and stream video – but not to render large numbers of continuous tone (i.e. raster)

images in real time. This required use of a graphics workstation, a high-end PC with a special graphics card or a game console with optimised graphics processors, as well as time-consuming software development. Although video games and 3D animation programs could render impressive numbers of pixels in real time, this was achieved by writing code that directly accesses hardware – something that very high-level media programming environments such as Processing and Flash/Flex could not do.

However, as the processing power and RAM size keep increasing, these differences between the graphics capacities of various hardware platforms and types of software are gradually disappearing. For example, *ImagePlot*,<sup>38</sup> a program which I wrote in 2009 using a high-level programming environment of imageJ (open source application for image processing commonly used in the sciences),<sup>39</sup> can render a 30,000 x 4,000 pixels image which shows 4535 *Time* covers in a few minutes on my Powerbook laptop (processor: 2.8 Ghz Intel Core 2 Duo; memory: 4GB 1067 Mhz DDR3). (Most of the time is spent on scaling down all the covers.) VisualSense<sup>40</sup> software that we developed in 2009–2010 with the National University of Singapore’s Multimodal Analysis Lab using Flash/Flex allows a user to define a number of graphs and change their positions and sizes. The graphs can use vector primitives (points, circles, rectangles) or they can show the actual images – thus allowing for interactive construction of direct visualisations. (Depending on the computer specifications, it can handle between 500 and 1000 images without slowing down.) Finally, the HiperView<sup>41</sup> application we developed (also in 2009) together with Calit2 Center of Graphics, Visualization and Virtual Reality (GRAVITY), takes advantages of the 286 megapixel resolution and significant memory of HIPerSpace to enable real-time interactive manipulation of image graphs which can contain up to 4000 images of *any* size.

I believe that direct visualisation methods will be particularly important for humanities, media studies and cultural institutions which now are just beginning to discover the use of visualisation but which eventually may adopt it as a basic tool for research, teaching and exhibition of cultural artefacts. (The first conference on visualisation in the humanities took place at MIT in May 2010.)<sup>42</sup> Those who study the humanities always focused on analysing and interpreting details of the cultural texts, be they poems, paintings, music compositions, architecture or, more recently, computer games, generative artworks and interactive environments. This is one of the key differences between humanities and sciences – at least, as they were practised until now. The

former are interested in particular artifacts (which can be taken to exemplify larger trends); those who study the latter are interested in general laws and models.

If scholars in the humanities start systematically using visualisation for research, teaching and public presentation of cultural artefacts and processes, the ability to show the artefacts in full detail is crucial. Displaying the actual visual media as opposed to representing them by graphical primitives helps the researcher to understand meaning and/or cause behind the pattern she may observe, as well as discover additional patterns.

While graphical reduction will continue to be used, this no longer constitutes the only possible method. The development of computers and the progress in their media capacities and programming environments now makes possible a new method for visualisation that I called 'direct visualisation' – that is, visualisation without reduction.<sup>43</sup>

## ACKNOWLEDGEMENT

The research on direct visualisation techniques and software by Software Studies Initiative was made possible by the generous support provided by California Institute for Telecommunication and Information (Calit2), University of California, San Diego's Center for Research in Computing and the Arts (CRCA), UCSD Chancellor's Office and the National Endowment for Humanities (NEH). The development of VisualSense software was made possible by a Mapping Asian Cultures grant from the University Research Council, National University of Singapore.

## NOTES

- [1] See [www.theusrus.de/Mondrian/](http://www.theusrus.de/Mondrian/).
- [2] For example: 'In contrast to scientific visualization, information visualization typically deals with nonnumeric, nonspatial, and high-dimensional data' (Chen 2005).
- [3] See [www.research.ibm.com/visual/projects/history\\_flow/](http://www.research.ibm.com/visual/projects/history_flow/).
- [4] See <http://www.aaronkoblin.com/work/flightpatterns/>.
- [5] See <http://processing.org/>.
- [6] See <http://prefuse.org/>.
- [7] See [http://britton.disted.camosun.bc.ca/beck\\_map.jpg](http://britton.disted.camosun.bc.ca/beck_map.jpg).
- [8] A number of definitions of information visualisation from the recent literature are available at [http://www.infovis-wiki.net/index.php?title=Information\\_Visualization](http://www.infovis-wiki.net/index.php?title=Information_Visualization).
- [9] See [www.math.yorku.ca/SCS/Gallery/milestone/sec5.html](http://www.math.yorku.ca/SCS/Gallery/milestone/sec5.html).
- [10] The historical data are from [www.math.yorku.ca/SCS/Gallery/milestone/sec4.html](http://www.math.yorku.ca/SCS/Gallery/milestone/sec4.html).
- [11] See <http://benfry.com/distellamap/>.
- [12] See <http://marumushi.com/projects/flickrgraph>.
- [13] See <http://ivl.slis.indiana.edu/research/>.
- [14] See <http://www.edwardtufte.com/tufte/minard>.
- [15] See [www.visualcomplexity.com/vc/project.cfm?id=696](http://www.visualcomplexity.com/vc/project.cfm?id=696).
- [16] See [www.google.com/trends](http://www.google.com/trends).
- [17] One important case which does not fit my analysis is the use already in the eighteenth century of different tones or colours to represent terrain elevation and relief in printed topographic maps. In these maps, tone or colour code qualitative data rather than categories.
- [18] See [http://en.wikipedia.org/wiki/Tag\\_cloud](http://en.wikipedia.org/wiki/Tag_cloud).
- [19] As an example, open source data visualisation software Mondrian 1.0 running on my 2009 Apple PowerBook laptop with 2.8 Ghz processor and 4 GB of RAM takes approximately seven seconds to render a scatter plot containing one million points.
- [20] Many additional examples of direct visualisation can be found in the field of motion graphics – film and television titles and graphics, commercials and music videos. In many motion graphics, text or images are animated to create dynamically changing meaningful patterns made from these media objects.
- [21] See <http://www.brendandawes.com/sketches/redux/>.
- [22] See <http://benfry.com/traces>.
- [23] I have created a few visualisations which show a whole book in a single image. See <http://www.flickr.com/photos/culturevis/sets/72157615900916808/>; <http://www.flickr.com/photos/culturevis/sets/72157622994317650/>. To display the whole text of Tolstoy's *Anna Karenina* in the smallest font which can be read, I had to create a 14,000 × 6,000 pixel image – well beyond the normal screen resolution today.
- [24] See <http://www.earstudio.com/projects/listeningpost.html>.
- [25] To see his taxonomy of network display methods, select 'filter by method' on [www.visualcomplexity.com/vc/](http://www.visualcomplexity.com/vc/).
- [26] See <http://en.wikipedia.org/wiki/Synecdoche>.
- [27] See <http://www.itsbeenreal.co.uk/index.php?/wwwwords/about-this-project/>.
- [28] See <http://www.turbulence.org/Works/song/>.
- [29] See <http://lab.softwarestudies.com/2008/09/cultural-analytics.html>.
- [30] See [http://vis.ucsd.edu/mediawiki/index.php/Research\\_Projects:\\_HIPerSpace](http://vis.ucsd.edu/mediawiki/index.php/Research_Projects:_HIPerSpace).
- [31] See [www.calit2.net](http://www.calit2.net).
- [32] See <http://www.flickr.com/photos/culturevis/4038907270/in/set-72157624959121129/>.
- [33] See <http://www.flickr.com/photos/culturevis/sets/72157623862293839/>.
- [34] See <http://www.flickr.com/photos/culturevis/4497385883/in/set-72157624959121129/>.

- [35] See <http://www.flickr.com/photos/culturevis/3951496507/in/set-72157622525012841/>.
- [36] A number of computer scientists have explored a related technique for browsing image collection where a part of a collection is displayed in a similar 'image graph' form. (For a summary of this work, see Marchand-Maillet and Bruno 2007, 5.) In most of the reported research, images are organised by visual similarity, which is calculated via computer image analysis. While this strategy is often useful for the analysis of cultural patterns, in many cases, such as the *Time* covers analysis, we want to see how visual features vary over time. Therefore we use original metadata (i.e. dates of publication) for one axis and measurement of one or more visual features (in this case, saturation) for the second axis.
- [37] The article is available at <http://www.manovich.net>.
- [38] See [www.flickr.com/photos/culturevis/sets/72157617847338031/](http://www.flickr.com/photos/culturevis/sets/72157617847338031/).
- [39] See <http://rsbweb.nih.gov/ij/>.
- [40] See [www.flickr.com/photos/culturevis/sets/72157623553747882/](http://www.flickr.com/photos/culturevis/sets/72157623553747882/).
- [41] See <http://lab.softwarestudies.com/2008/09/cultural-analytics.html>.
- [42] See [hyperstudio.mit.edu/h-digital/](http://hyperstudio.mit.edu/h-digital/).
- [43] It is possible, however, that our interactive interfaces to visualisations are effective precisely because they do provide certain reduction functions. I am thinking in particular about the zoom command. We zoom into direct visualisation such as that of *Time* covers to examine the details of particular covers. We zoom out to see the overall trends. When we do that, the images are gradually reduced in size, eventually becoming small colour dots.
- Chen, C. 2005. Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications* 25 (4) (July–August): 12–16.
- Friendly, Michael, and Daniel J. Denis. 2010. Milestones in the history of thematic cartography, statistical graphics, and data visualization, section 5. <http://www.math.yorku.ca/SCS/Gallery/milestone/sec5.html>.
- Keim, D.A., F. Mansmann, J. Schneidewind, and H. Ziegler. 2006. Challenges in visual data analysis. *Proceedings of Information Visualization (IEEE) IV* (2006): 9–16.
- Latour, Bruno. 2009. Tarde's idea of quantification. In *The social after Gabriel Tarde: Debates and assessments*, edited by Mattei Candea. <http://www.bruno-latour.fr/articles/article/116-TARDE-CANDEA.pdf>.
- Manovich, Lev. 2002. Data visualization as new abstraction and anti-sublime. <http://manovich.net/articles/>.
- Marchand-Maillet, S., and E. Bruno. 2007. State of the art image collection overviews and browsing. <http://www.multimatch.org/docs/publicdels/D1.1.2.pdf>.
- Purchase, H.C., N. Andrienko, T.J. Jankun-Kelly, and M. Ward. 2008. Theoretical foundations of information visualization. In *Information visualization: Human-centered issues and perspectives*, edited by A. Kerren, J.T. Stasko, J. Fekete and C. North. *Lecture notes in computer science*, vol. 4950, 46–64. Berlin, Heidelberg: Springer-Verlag.
- Tufte, Edward. 1983. *The visual display of quantitative information*. Cheshire, CT: Graphics Press.
- . 1990. *Envisioning information*. Cheshire, CT: Graphics Press.
- . 1997. *Visual explanations: Images and quantities, evidence and narrative*. Cheshire, CT: Graphics Press.
- . 2006. *Beautiful evidence*. Cheshire, CT: Graphics Press.
- van Ham, Frank, Martin Wattenberg, and Fernanda B. Viégas. 2009. Mapping text with phrase nets. *IEEE Transactions on Visualization and Computer Graphics* 15 (6): 1169–76.
- Venturi, Robert, Stefan Izenour, and Denise Scott Brown. 1972. *Learning from Las Vegas*. Cambridge, MA: MIT Press.

## REFERENCES

Ball, Philip. 2004. *Critical mass*. London: Arrow Books.