# What is Visualization?[1]

## © Lev Manovich 2010

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 answer 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. [William Playfair, 1805.[2]]

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? [Robert Venturi, Stefan Izenour, Denise Scott Brown, 1972.[3]]

'Whole' is now nothing more than a provisional visualization which an 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. [Bruno Latour, 2009.[4]]

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, 2008.[5]]

Visualization is ready to be a mass medium. [Fernanda B. Viégas and Martin Wattenberg, 2010.[6]]

2010. Museum of Modern Art in New York presents a dynamic visualization of its collection on 5 screens created by Imaginary Forces. New York Times regularly features custom visualizations both in its print and web editions created by the in-house The NYTimes interactive team. The web is crawling with numerous sophisticated visualization 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 this data. If you

want to visualize our own data set, Many Eyes, Tableau Public and other sites offer free visualization tools. 300 years after William Playfair's amazement at the cognitive power of information visualization, it looks like that finally many others are finally getting it.

What is information visualization? Despite the growing popularity of infovis (a common abbreviation for "information visualization"), 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 would clearly separate it from other related fields such as scientific visualization and information design. So lets start with a provisional definition that we can modify later. Lets define information visualization 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 a 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 visualization – such as the distinctions between static, dynamic (i.e. animated) and interactive visualization – 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 the examples of such definitions: "Information visualization (INFOVIS) is the communication of abstract data through the use of interactive visual interfaces."[7] "Information visualization utilizes computer graphics and interaction to assist humans in solving problems."[8]

Interactive graphic interfaces in general, and interactive visualization 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 OS to multiple coordinated views available in some visualization software such as Mondrian.[9] However, regardless of whether you are looking at a visualization printed on paper or a dynamic arrangement of graphic elements on your computer screen which you generated using interactive software and which you 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 visualization is distinct from scientific visualization in that the latter uses numerical data while the former uses non-numeric data such as text and networks of relations.[10] 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 visualization 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 to it.[11] 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 US to create an animated map which display the pattern formed by their movement over a 24-hour period.[12]

Rather than trying to separate information visualization and scientific visualization using some a priori idea, let's instead enter each phrase in Google image search and compare the results. Most images returned by searching for "information visualization" are two dimensional and use vector graphics - points, lines, curves, and other simple geometric shapes. The majority of images returned when searching for "scientific visualization" 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 necessary use different types of data but because they privilege different visual techniques and technologies.

Scientific visualization and information visualization come from different cultures (science and design); their development corresponds to different areas of computer graphics technology. Scientific visualization developed in the 1980s along with the field of 3D computer graphics, which at that time required specialized graphics workstations. Information visualization developed in the 1990s along with the rise of desktop 2D graphics software and the adoption of PCs by designers; its popularity accelerated in 2000s – the two key factors being the easy availability of big data sets via APIs provided by major social network services since 2005 and new high level programming languages specifically designed for graphics (i.e., Processing[13]) and software libraries for visualization (for instance, Prefuse[14]).

Can we differentiate information visualization from information design? This is more tricky, but here is my way of doing it. Information design starts with the data that already has 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.[15] In contrast, the goal of information visualization is to discover the structure of a (typically large) data set. This structure is not known a priori; a

visualization is successful if it reveals this structure. A different way to express this is to say that information design works with information, while information visualization works with data. As it 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 visualization 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 19[th] and 20[th] century that are known to many via the examples collected in the pioneering books by Edward Tufte?[16] Do they constitute infovis as we understand it today? As I already noted, most definitions by the researchers working within Computer Science tend to equate information visualization with the use of interactive computer graphics.[17] Using software, we can visualize much larger data sets than it was possible previously; create animated visualization; show how processes unfold in time; and, most importantly, manipulate visualizations 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 visualization - mapping some properties of the data into a visual representation. Similarly, while availability of computers led to the development of new visualization techniques (scatter plot matrix, treemaps, etc.), the basic visual language of infovis remained the same as it was in the 19[th] 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 visualization.

Reduction and Space

In my view, the practice of information visualization from its beginnings in the second part of the 18<sup>th</sup> century until today relied on two key principles. The first principle is REDUCTION. Infovis uses graphical primitives such as points, strait 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 schematization. We throw away 99% of what is specific about each object to represent only 1%- in the hope of revealing patterns across this 1% of objects' characteristics.

Information visualization 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 19<sup>th</sup> century when in the course of just a few decades almost all graph types commonly found today in statistical and charting programs are invented.[18] This development of the new techniques for visual reduction parallels the reductionist trajectory of modern science in the 19<sup>th</sup> century. Physics, chemistry, biology, linguistics, psychology and sociology propose that both natural and 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 behavior emerge out of interaction of SIMPLE ELEMENTS.

Even more direct is the link between 19[th] century infovis and the rise of social statistics. Philip Ball summarizes 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'.[19]

In the first part of the 19[th] century many scholars inluding Adolphe Quetelet, Florence Nightingale, Thomas Buckle, and Francis Galton used statistics to look for "laws of society." This inevitably involved summarization 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 standard today were invented during this time for the purposes of representations of such summarized data. According to Michael Friendly and Daniel J. Denis, between 1800 and 1850, "In 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."[20]

Do all these different visualization 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 practiced 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, OpenOffice, etc.[21]

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 into topology and geometry. Other less important properties of the objects are represented through different visual dimensions  - tones, shading patterns, colors, 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  published in 1786 and became commonplace in the early 19th 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; etc.

Lets take another example - network visualizations which function today as distinct symbols of "network society" (see Manuel Lima's authoritative gallery visualcomplexity.com which currently houses over 700 network visualization projects). Like bar charts and line graphs, network visualizations also privilege spatial dimensions: position, size, and shape. Their key addition is the use of strait or curved lines to show connections between data objects. For example, in DISTELLAMAP (2005) Ben Fry connects pieces of code and data by lines to show the

dynamics of the software execution in Atari 2600 games.[22] In Marcos Weskamp's Flickr

Graph (2005) the lines visualize the social relationships between users of flickr.com.[23] (Of

course, many other visual techniques can also be used to 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.[24]

I believe that the majority of information visualization practice from the second part of

the 18th 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 visualizations ranging from famous

dense graphic showing Napoleon's March on Moscow by Charles Joseph Minard (1869)[25] to the

recent The Evolution of The Origin of Species by Stefanie Posavec and Greg

McInerny (2009).[26] 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 visualizations use colors, 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 use line graphs to compare search volumes

for different words or phrases; each line is rendered in a different color.[27] However the same

visualization could have simply used labels attached to the lines - without different colors. In this

case, color ads readability but it does not add new information to the visualization.

The privileging of spatial over other visual dimensions was also true of plastic arts in

Europe between 16th and 19th 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 color was 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 colors, textures and other visual parameters.  In classical Asian "ink and wash painting" which first appeared in 7[th] century in China and was later introduced to Korea and then Japan (14[th] century), color did not even appeared. 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 visualizations where the main dimension is color –for instance, a common traffic light which "visualizes" the three possible behaviors of a car driver: stop, get ready, go. This example shows that if we fix spatial parameters of visualization, color 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 visualization have the same values on spatial dimensions, our visual system can focus on the differences represented by colors, or other non-spatial variables.

Why do visualization designers – be they the inventors of graph and chart techniques at the end of the 18[th] and early 19[th] century, or millions of people who now use these graph types in their reports and presentations, or the authors of more experimental visualizations featured on infoaesthetics.com and visualcomplexity.com - privilege spatial variables over other kinds of visual mappings? In other words, why color, tone, transparency, and symbols are used to represent secondary aspects of data while the spatial variables are reserved for the most important dimensions? Without going into the details 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 visualizations 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 segments 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 color (think of twentieth century art) – but in a everyday perception, spatial properties are what matters most. How close are two people to each other; the expression on their faces; their relative size which allows the observer to estimate their distance from her; the characteristic shapes of different objects which allows her to recognize 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 maybe the reason why all standard techniques for making graphs and charts developed in the $18^{th}$ – $20^{th}$ centuries 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 visualizations on computer monitors, color become the norm.  Color printing is still significantly more expensive than using a single color – so even

today science journals are printed in black and white. Thus, the extra cost associated with creating and printing color graphics during the last two centuries was probably an important factor responsible for privileging of spatial variables.

When color, shading, and other non-spatial visual parameters were used in visualizations created in the 19[th] and most of the 20[th] century, they usually represented only small number of discrete values – i.e. they acted as "categorical variables." However today the fields of computer-based scientific visualization, geo-visualization, 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, computers monitors can show 16 million unique colors. Therefore color, shading and transparency are now commonly employed in these fields to show continuously varying qualities such temperature, gas density, elevation, gravity waves, etc. Does not this contradict my statement that spatial arrangement is key to information visualization?

We can solve this puzzle if we take into account a fundamental difference between information visualization and scientific visualization / geovisualization, which I did not yet mention. Infovis uses ARBITRARY spatial arrangements of elements to represent the relationships between data objects. Scientific, medical and geovisualization typically work with A PRIORI FIXED SPATIAL LAYOUT OF THE REAL PHYSICAL OBJECTS such as a brain, a coastline, a galaxy, etc. Since the layout in such visualizations is already fixed and can't be arbitrary manipulated, color and/or other non-spatial parameters are used instead to show new information. A typical example of this strategy is a *heat map* which use color hue and saturation to overlay information over a spatial map.[28]

The two key principles that I suggested  – *data reduction* and *privileging of spatial variables* - do not account for all possible visualizations produced during last 300 years. However, they are sufficient to separate infovis (at least as it was commonly practiced 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.

## Visualization Without Reduction

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

If we survey the practice of infovis from the 18<sup>th</sup> until the end of the 20<sup>th</sup> century, the idea that visualization takes data that is not visual and maps it into a visual domain indeed works quite well. However, it seems to longer adequately describe certain new visualization techniques and projects developed since the middle of the 1990s. Although these techniques and projects are commonly discussed as "information visualization," 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 visualization method for which we don't yet have an adequate name.

Consider a technique called tag cloud.[29] This technique was popularized 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 VISUALIZATION: creating new visual representations from the actual visual media objects, or their parts. Rather than representing text, images, video or other media though new visual signs such as points or rectangles, media visualizations build new representations out of the original media. Images remain images; text remains text.

In view of our discussion of data reduction principle, we can also call this method DIRECT VISUALIZATION, or VISUALIZATION WITHOUT REDUCTION. In direct visualization, the data is reorganized into a new visual representation that preserves its original form. Usually, this does involve some data transformation such as changing data size. For instance, text 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 don't substitute media objects by 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 "visualization without reduction" refers to this preservation of a much richer set of properties of data objects when we create visualizations directly from them.

Not all direct visualization techniques such as a tag cloud originated in the 21$^{st}$ 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 visualization 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 visualization method, 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 into its font size was counter-intuitive – so it was not invented. In contrast, tag cloud technique is a typical expression of what we can call "software thinking" – i.e. 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, etc. If we take these two capacities for granted, the idea to arbitrary change the size of words based on some information - such as their frequency in a text - is something we may expect to be "actualized" in the process of cultural evolution. (In fact, all contemporary interactive visualization techniques rely on the same two fundamental capacities.)

The rapid growth in the number and variety of visualization projects, software applications, and web services since the late 1990s was enabled by the advances in computer graphics capacities of PCs including both hardware (processors, RAM, displays) and software (C and Java graphics libraries, Flash, Processing, Flex, Prefuse, etc.) These developments both popularized information visualization and also fundamentally changed its identity by foregrounding animation, interactivity and also more complex visualizations that represent connections between many more objects than previously.[30] But along with these three highly

visible trends, the same advances also made possible "direct visualization" approach – although it has not been given its own name so far.

### Direct Visualization: Examples

Lets discuss three well-known projects which exemplify "direct visualization": Listening Post, *Cinema Redux*, and *Preservation* of Selected Traces.[31]

Cinema Redux was created by interactive designer Brendan Dawes in 2004.[32] Dawes wrote a program in Processing that sampled a film at the rate of one frame per second and scaled each frame to 8x6 pixels. The program then arranged these minuate frames in a rectangular grid with every row representing a single minute of the film. Although Dawes could have easily continue this process of sampling and remapping – for instance, representing each frame though its dominant color - he chose instead to use the actual scaled down frames from the film. The resulting visualization represents a trade-off between the two possible extremes: preserving all the details of the original artifact and abstracting its structure completely. Higher degree of abstraction may make the patterns in cinematography and narrative more visible but it would also remove the viewer further from the experience of the film. Staying closer to the original artifact 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 are not the particular parameters which Dawes used for Cinema Redux but that he reinterpreted the previous constant of visualization practice as a variable. Previously infovis designers mapped data into new diagrammatic representation consisting from 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 chose 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 – i.e. expanding typical graphical primitives (points, rectangles, etc.) into the actual data objects (film frames).

Before software, visualization 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 artifacts without quantifying them. As demonstrated by *Cinema Redux*, these manipulations can successfully make visible the relations between a large number of these artifacts. Of course, such visualization without quantification is made possible by the a priori quantification required to turn any analog data into a digital representation. In other words, it is the "reduction" first performed by the digitization process which paradoxically now allows us to visualize the patterns across sets of analog artifacts without reducing them to graphical signs.

For another example of direct visualization, let's turn to Ben Fry's *Preservation of Selected Traces* (2009).[33] This web project is an interactive animation of the complete text of Darwin's Evolution of the Species. Fry uses different colors 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 text from edition to edition, with sentences and passages deleted, inserted and re-written. In contrast to typical animated information visualizations 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 book's 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 compete 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.[34] Instead sentences are rendered as tiny rectangles in different colors. 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 visualization.

Let's add one more example – LISTENING POST by Ben Rubin and Mark Hansen (2001).[35] 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 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 visualization 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 is not – it is 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 always differs, 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 representative of direct visualization method: the patterns it presents depend as much on what all the text fragments that appear on a 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 identifies what he calls a "syntax" of network visualizations – commonly used layouts such as radial convergence, arc diagrams, radial centralized networks, and others.[36] The key difference between most of these network visualizations and LISTENING POST lies in the fact that the former often rely on the existing algorithms to create visualization layouts. 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 layout's 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. Bruno Latour's article (quoted in the beginning) argues that our ability to create "a provisional visualization which can be modified and reversed" allows us to think differently since any "whole" we can construct now is just one of numerous others.  But a full eight years

before Latour wrote that essay, Listening Post already beautifully staged this new epistemological paradigm enabled by interactive visualization.

The three influential projects I have considered demonstrate that in order to highlight patterns in the data we don't have to reduce it by representing data objects via abstract graphical elements. We also don't have to summarize the data as it 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 visualization" an image has to show 100% of the original data – every word in a text, every frame in a movie, etc. Out of the three examples I just discussed, only Preservation of Selected Traces does  so.  Both Cinema Redux and Listening Post do not use all the available data – instead they sample it. The first project samples a feature film at the fixed rate of 1 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 visualizations are not the result of remapping of the data into some new representational 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.[37] (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% of the data.  But you already have been employing this strategy - if you ever used a magic marker to highlight important passages of a printed text. Although text highlighting normally is not thought as visualization, we can see that in fact it is an example of "direct visualization without sampling."

Cinema Redux and Preservation of Selected Traces also break away from the second key principle of traditional visualization - 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 visualize is not the same as the typical data used in infovis. Films or books are not just a collection of data objects; they are narratives made from these objects (i.e. the data has a sequential order). Although it is certainly possible to create effective visualizations that remap a narrative sequence into a completely new spatial structure as in Listening Post (see also Writing Without Words by Stefanie Posavec[38] and The Shape of Song by Martin Wattenberg[39]), 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 temporal 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, meaning, and its experience by the users. While we consciously or unconsciously notice many of these patterns during watching / reading / interacting with the work, projecting time into space - laying out movie frames, book sentences, magazine pages in a single image  - gives us new possibilities to study these artifacts. Thus, *space* plays a crucial role in direct visualization 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 visualization created at my lab, Software Studies Initiative (softwarestudies.com).[40] Inspired by the artistic projects which pioneered the direct visualization approach as well by the resolution and real-time capabilities of SUPERVISUALIZATION interactive systems such as HIPerSpace (35,840 by 8,000 pixels, 286,720,000 pixels total[41]) developed at California Institute for Telecommunication and Information (Calit2)[42] 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 visualizations 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 it. 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).[43] We also compared the use of images in SCIENCE and POPULAR SCIENCE magazines by visualizing approximately 10,000 pages from each magazine during first decades of their publication (*The Shape of Science*, William Huber, Lev Manovich, Tara Zapel, 2010).[44] Our most data-intensive direct visualization is 44,000 by 44,000 pixels; it shows 1,074,790 Manga pages organized by their stylistic properties (*Manga Style Space*, Lev Manovich and Jeremy Douglass, 2010).[45]

Like CINEMA REDUX, *Mapping* TIME and THE SHAPE OF SCIENCE make equal the values of spatial variables to reveal the patterns in the content, colors, 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 visualizations use only one dimension, the sequence of images wrapped around into a number of rows to make it easier to see patterns without having to visually scan a very long image. However, we can turn such one-dimensional

image timelines into 2D, with the second dimension communicating additional information. Consider a 2D timeline of Tɪᴍᴇ Mᴀɢᴀᴢɪɴᴇ covers we created (*Timeline*, Jeremy Douglass and Lev Manovich, 2009).[46] 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 colors) 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 color saturation gradually increases as Tɪᴍᴇ publication reached 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 the Tɪᴍᴇ covers gradually rises 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 21$^{st}$ century, this trend is reversed: the covers now have less contrast and less saturation.

The strategy used in this visualization is based on a familiar technique, a scatter graph. However, if a normal scatter graph reduces the data displaying each object to a point, we display the data in its 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."[47]


## Wʜᴀᴛ ɪs Vɪsᴜᴀʟɪᴢᴀᴛɪᴏɴ?

In an article on the emerging practice of artistic visualization written in 2002, I defined visualization as "a transformation of quantified data which is not visual into a visual representation." At that time I wanted to stress that visualization participates in the reduction projects of modern science and modern art which led to the choice of the article's title:  "Data

Visualization as New Abstraction and Anti-Sublime."[48] I think that this emphasis was

appropriate given the types of infovis typically created at that time. (Although I used 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 visualization 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 visualized is already visual – text, film frames,

magazine covers. In other words, these projects create new visual representations our of the

original visual data without translating it into graphic signs. They also often break away from the

second key principle of infovis  - mapping of most important data dimensions into spatial

variables.

So is "direct visualization" actually constitutes 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.  Alternatively, we can revise our understanding of what infovis is.

Given that all direct visualizations we looked at aim to make visible patterns and relations

in the data, this aim certainly aligns them with infovis as it developed during last 300 years. It is

also relevant to note that some of the of most well-known infovis projects of the last 15 years

follow direct visualization approach. This is true of Cinema Redux and Preservation of

Selected Traces and other seminal projects which I did not discussed in detail such as the

*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

from vector elements but from 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 IEEE InfoVis 2009 conference also operates within a direct visualization paradigm.[49]

Does this mean that what we took to be the core principle of information visualization during its first three centuries – reduction to graphic primitives – was only a particular historical choice, an artifact 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 - color, transparency, texture, and any other visual parameter of any part of an image allows us to start using these non-spatial parameters to represent the key dimensions of the data. This is already common in scientific, medical and geo-visualization – but not yet in information visualization.

Why has infovis continued to rely on computer-generated VECTOR graphics during 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 visualization. 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 a use of graphics workstation, a high-end PC with a special graphics card or a game console with optimized 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 software are gradually disappearing. For example, the ImagePlot[50] program which I wrote in 2009 using high-level programming environment of imageJ (open source application for image processing commonly used in the sciences[51]) can render a 30000x4000 pixels image which shows 4535 Tɪᴍᴇ 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[52] software that we developed in 2009-2010 with 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 visualizations. (Depending on the computer specifications, it can handle between 500 and 1000 images without slowing down.) Finally, the HiperView[53] 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 visualization methods will be particularly important for humanities, media studies and cultural institutions which now are just beginning to discover the use of visualization but which eventually may adopt it as a basic tool for research, teaching and exhibition of cultural artifacts. (The first conference on visualization in the humanities took place at MIT in May 2010[54]). Humanists always focused on analyzing 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 practiced until now. The former are interested in particular artifacts (which can be taken to exemplify larger trends); the latter are interested in general laws and models.

If humanists start systematically using visualization for research, teaching and public presentation of cultural artifacts and processes, the ability to show the artifacts in full detail is crucial. Displaying the actual visual media as opposed to representing it 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 is no longer 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 visualization that I called "direct visualization" – i.e., visualization without reduction.[55]

Lev Manovich
March - October 2010

# **Notes**

[1] In order to follow the links in this document automatically, you will need download the pdf file from the main table of contents and open the file locally, on your device.  This necessity is due to problem with our OJS software; we apologize.

[2] William Playfair, An Inquiry into the Permanent Causes of the Decline and Fall of Powerful and Wealthy Nations [1805]; 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."

http://books.google.com/books?id=MLvIAAAAMAAJ

[3] Robert Venturi, Stefan Izenour, Denise Scott Brown], Learning from Las Vegas (1972). Emphasis is in the original:

http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=3723&mode=toc

[4] Bruno Latour, Tarde's Idea of Quantification, in The Social After Gabriel Tarde: Debates and Assessments, ed. Mattei Candea (2009): http://www.bruno-latour.fr/articles/

[5] Eric Rodenbeck (Stamen Design), keynote lecture at Emerging Technology 2008, March 4, 2008: http://en.oreilly.com/et2008/public/schedule/detail/1588

[6] Fernanda B. Viégas and Martin Wattenberg, an interview for infosthetics.com (May 2010): http://infosthetics.com/archives/2010/05/interview_fernanda_viegas_and_martin_wattenberg_from_flowing_media.html

[7] Keim, D.A.; Mansmann, F. and Schneidewind, J. and Ziegler, H.. "Challenges in Visual Data Analysis", Proceedings of Information Visualization (IV 2006), IEEE, p. 9-16, 2006.

[8] Purchase, H. C., Andrienko, N., Jankun-Kelly, T. J., and Ward, M. 2008. "Theoretical Foundations of Information Visualization", Information Visualization: Human-

Centered Issues and Perspectives, A. Kerren, J. T. Stasko, J. Fekete, and C. North, Eds. Lecture Notes In Computer Science, vol. 4950. Springer-Verlag, Berlin, Heidelberg, 46-64.

9 http://www.theusrus.de/Mondrian/.

10 For example: "In contrast to scientific visualization, information visualization typically deals with nonnumeric, nonspatial, and high-dimensional data." Chen, C. "Top 10 Unsolved Information Visualization Problems", IEEE Computer Graphics and Applications, 25(4):12-16, July-Aug. 2005.

11 http://www.research.ibm.com/visual/projects/history_flow/.

12 http://www.aaronkoblin.com/work/flightpatterns/.

13 http://processing.org/.

14 http://prefuse.org/.

15 http://britton.disted.camosun.bc.ca/beck_map.jpg.

16 Edward Tufte, The Visual Display of Quantitative Information. Cheshire, CT: Graphics Press, 1983. Edward Tufte, Envisioning Information. Cheshire, CT: Graphics Press, 1990. Edward Tufte, Visual Explanations: Images and Quantities, Evidence and Narrative. Cheshire, CT: Graphics Press, 1997. Edward Tufte, Beautiful Evidence. Cheshire, CT: Graphics Press, 2006.

17 A number of definitions of information visualization from the recent literature is available at http://www.infovis-wiki.net/index.php?title=Information_Visualization.

18 http://www.math.yorku.ca/SCS/Gallery/milestone/sec5.html.

19 Philip Ball, Critical Mass. London: Arrow Books, 2004. Pp. 64-65.

[20] Michael Friendly and Daniel J. Denis, MILESTONES IN THE HISTORY OF THEMATIC CARTOGRAPHY, STATISTICAL GRAPHICS, AND DATA VISUALIZATION, section 5 < http://www.math.yorku.ca/SCS/Gallery/milestone/sec5.html>.

[21] The historical data is from www.math.yorku.ca/SCS/Gallery/milestone/sec4.html.

[22] http://benfry.com/distellamap/.

[23] http://marumushi.com/projects/flickrgraph.

[24] http://ivl.slis.indiana.edu/research/.

[25] http://www.edwardtufte.com/tufte/minard.

[26] http://www.visualcomplexity.com/vc/project.cfm?id=696.

[27] http://www.google.com/trends.

[28] One important case which does not fit my analysis is the use of different tones or colors to represent terrain elevation and relief in printed topographic maps already in the $18^{th}$ century. In these maps, tone or color codes qualitative data rather than categories.

[29] http://en.wikipedia.org/wiki/Tag_cloud.

[30] As an example, open source data visualization software Mondrian 1.0 running on my 2009 Apple PowerBook laptop with 2.8 Ghz processor and 4 GB of RAM takes approximately 7 seconds to render a scatter plot containing 1 million points.

[31] Many additional examples of direct visualization can be found in the field of motion graphics - film and TV 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.

[32] http://www.brendandawes.com/sketches/redux/

[33] http://benfry.com/traces/.

[34] I have created a few visualizations 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 a smallest font which can be read, I had to make 14,000 x 6,000

pixels – well beoynd the normal screen resolution today.

[35] http://www.earstudio.com/projects/listeningpost.html.

[36] To see his taxonomy of network display methods, select "filter by method" on

www.visualcomplexity.com/vc/.

[37] http://en.wikipedia.org/wiki/Synecdoche.

[38] http://www.itsbeenreal.co.uk/index.php?/wwwords/about-this-project/.

[39] http://www.turbulence.org/Works/song/.

[40] http://lab.softwarestudies.com/2008/09/cultural-analytics.html.

[41] http://vis.ucsd.edu/mediawiki/index.php/Research_Projects:_HIPerSpace.

[42] http://www.calit2.net.

[43] http://www.flickr.com/photos/culturevis/4038907270/in/set-72157624959121129/.

[44] http://www.flickr.com/photos/culturevis/sets/72157623862293839/.

[45] http://www.flickr.com/photos/culturevis/4497385883/in/set-72157624959121129/.

[46] http://www.flickr.com/photos/culturevis/3951496507/in/set-72157622525012841/.

[47] 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 S. Marchand-Maillet, E. Bruno, State of the Art Image Collection

Overviews and Browsing (2006), p. 5. <www.multimatch.org/docs/publicdels/D1.1.2.pdf>. In

most of the reported research, images are organized 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 *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.

[48] The article is available at www.manovich.net.

[49] Frank van Ham, Martin Wattenberg, Fernanda B. Viégas, Mapping Text with Phrase Nets, IEEE InfoVis 2009.

[50] http://www.flickr.com/photos/culturevis/sets/72157617847338031/.

[51] http://rsbweb.nih.gov/ij/.

[52] www.flickr.com/photos/culturevis/sets/72157623553747882/.

[53] http://lab.softwarestudies.com/2008/09/cultural-analytics.html.

[54] hyperstudio.mit.edu/h-digital/.

[55] It is possible however that our interactive interfaces to visualizations are effective precisely because they do provide certain reduction functions. I am thinking in particular about zoom command. We zoom into direct visualization such as covers of Time Magazine 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 color dots.