Visualization Methods for Media Studies

Lev Manovich

In the first decade of the 21[st] century, the researchers in the humanities and humanistic social sciences have gradually started to adopt computational and visualization tools. The majority of this work often referred as "digital humanities" has focused on textual data (e.g., literature, historical records, or social media) and spatial data (e.g., locations of people, places, or events).[1] However, visual media have remained outside of the new computational paradigm. To fill this void, in 2007 I established the Software Studies Initiative at University of California, San Diego.[2] Our first goal was to develop easy to use techniques for visualization and computational analysis of large collections of images and video suitable for researchers in media studies, the humanities, and the social sciences who do not have technical background, and to apply these techniques to progressively large media data sets. Our second goal was theoretical - to examine existing practices and assumptions of visualization and computational data analysis (thus the name "Software Studies"), and articulate new research questions enabled by humanistic computational work with "big cultural data" in general, and visual media specifically.[3]

This chapter draws on the number of my articles written since we started the lab where I discuss history of visualization, the techniques that we developed for visualizing large sets of visual media, and their applications to various types of media.[4] The reader is advised to consult these articles on the details of visualization methods presented and detailed analysis of their applications. The first purpose of this chapter is to bring together the key theoretical points developed across these articles.

In doing this, I also want to articulate the connections between some of the key concepts involved in visualizing media for humanities research - "artifact," "data," "metadata," "feature",

---

[1] For recent discussions of digital humanities, see David M. Berry, ed., *Understanding Digital Humanities* (Palgrave Macmillan, 2012); Matthew K. Gold, ed. *Debates in the Digital Humanities* (University Of Minnesota Press, 2012); Katherine Hayles, *How We Think: Digital Media and Contemporary Technogenesis* (University Of Chicago Press, 2012); Anne Burdick, Johanna Drucker, Peter Lunenfeld, Todd Presner, Jeffrey Schnapp, Digital *Humanities* (The MIT Press, 2012); Stephen Ramsay, *Reading Machines: Toward an Algorithmic Criticism* (University of Illinois Press, 2011).

[2] www.softwarestudies.com.

[3] To separate this research from many other kinds of work included in 2000s under the umbrella term "digital humanities," I introduced the term Cultural Analytics to refer to the use of visualization and quantitative analysis of large sets of visual and interactive artifacts for humanities research and teaching. See Lev Manovich, "Cultural Analytics: Visualizing Cultural Patterns in the Era of 'More Media'," Domus (Milan), 2009.

[4] The key articles are: Lev Manovich, "What is visualization?" *Visual Studies*, vol. 26, no.1 (2011): 36-49; Lev Manovich, "How to compare one million images?" in *Understanding Digital Humanities*, ed. David Berry (New York: Palgrave Macmillan, 2012); Lev Manovich, "Media Visualization: Visual Techniques for Exploring Large Media Collections," in *Media Studies Futures*, ed. Kelly Gates (Blackwell, 2012).

"mapping," and "remapping." We can relate these concepts in three ways. Firstly, we can look at these and other related concepts as series of oppositions: artifact vs. data, data vs. metadata, close reading vs. distant reading. Secondly, since the combination of these concepts correspond to fundamental conceptual steps used in various visualization methods, we can examine theoretically at each of these steps (translating from artifacts to data, adding new metadata, extracting features, mapping and remapping from data to a visual representation.)

Thirdly, we can organize our discussion in terms of these methods. For example, visualization can show the metadata about the artifacts or the actual artifacts; a researcher can use existing metadata or add new ones. The conceptual characterization of these fundamental methods is the third goal of this chapter. It organizes the methods along two conceptual dimensions. The first dimension describes what is the prime object being visualized - data or metadata. The second dimension describes the two key ways of augmenting the original data with new information used in visualization – manual annotation or automatic feature extraction.

Since my lab focused on working with visual media data sets - photography, images of art, films, cartoons, motion graphics, video games, book pages, magazine covers and pages, and so on – all the methods described will be immediately applicable to all types of visual media. However, as I will explain, not all of them will work with other types of media because of the particular properties of images and human vision.

Artistic Visualization and Humanities

[INSERT FIGURE 1 HERE]

There is a multitude of visualization techniques available today.[5] The systematic history of their development, the connections to the need of modern societies and science analyze and manage progressively larger amounts of data, and, more recently, the increasing capacities of computer technologies, remains to be written, but at least key milestones are known.[6] For example, popular software such as Excel, Tableau, manyeyes, and others offer a set of graphing techniques which were developed already in the first decades of the 19th century - pie charts, bar charts, scatterplots, radar charts, histograms, etc. The same period also witnessed the development of 2D thematic maps that visualized data on variety of topics. The adoption of computers led to many new techniques, as well as gradual increase in information density of representations since software programs could visualize much larger amounts of data when it was practical to do by hand. The rapid development of 3D computer graphics technologies in the 1980s made possible the development of the new field of scientific visualization. The next

---

[5] See Nathan Yau, *Visualize This: The FlowingData Guide to Design, Visualization, and Statistics* (Wiley, 2011).
[6] Michael Friendly and Daniel J. Denis, *Milestones in the History of Thematic Cartography, Statistical Graphics, and Data Visualization*, http://datavis.ca/milestones/; http://www.datavis.ca/gallery/index.php.

wave was the rise of information visualization in the 1990s that introduced new 2D techniques (such as hyperbolic trees and treemaps) for representing non-numerical data.

In the late 1990s, information visualization started to attract attention of new media artists; by 2004 multitude of projects they created reached a point where it became meaningful to talk about the new area of "artistic visualization."[7] Although this new area of culture continued to grow, with visualization projects included in major museum exhibitions, the term itself remained problematic. (For one thing most celebrated examples of artistic visualizations were created by professionally trained designers Ben Fry and Lee Byron, and scientist Martin Wattenberg). One way to define artistic visualization is by contrasting it to the "normal" use of visualization in science, business and mass media. If these fields use visualization functionally, with a designer aiming to represent the relationships in a data given to her by the client without making any independent statement about it (we call this position "design neutrality"), artistic visualization projects deliberately aim to make such statements. The goal, in other words, is not a representation of data for its own sake but rather a statement about the world and human beings made through particular choices of the data sets and their presentation.[8]

As artistic visualization became popular with digital artists and designers, the number of people doing this work kept increasing. (Significant factors here were the development of Processing high-level graphics language designed specifically for artists, and the availability of data from major social media sites via their APIs.) A constant competition on the level of form became another distinguishing feature of artistic visualization. We can say that the history of visualization entered a new "modernist" stage where the invention of new techniques (or, at least, new variations of the existing techniques) came to be valued for its own sake. Indeed, a survey of the most influential artistic visualization projects of 2000s shows that none of them used already well-know visualization techniques but instead defined new ones. Some of these new techniques were given explicit names and since their introduction in particular projects were adopted by other designers (for example, arc diagrams from *The Shape of Song* by Martin Wattenberg, 2001, Streamgraph from Lee Byron's *Listening History*, 2006); others only appeared in unique visualization projects (Fernanda B. Viégas and Martin Wattenberg, *History Flow*, 2002; *On the Origin of Species: The Preservation of Favoured Traces* by Ben Fry, 2009.)

However, the artistic projects that are able to introduce really new visualization techniques are exceptions; the majority of projects are only able to distinguish themselves by customizing already existing techniques. For example, consider visualcomplexity.com, the influential

---

[7] See Lev Manovich, "Data Visualization as New Abstraction and Anti-Sublime," *SMAC! 3* (San Francisco, 2002), http://lab.softwarestudies.com/2008/09/cultural-analytics.html; Andrew Vande Moere, "About the Information Aesthetics Weblog" (12/2004), http://infosthetics.com/information_aesthetics_about.html; Fernanda B. Viégas and Martin Wattenberg, "Artistic Data Visualization: Beyond Visual Analytics," *Proceedings of the 2nd International Conference on Online Communities and Social Computing*. Springer-Verlag Berlin, 2007, http://www.research.ibm.com/visual/papers/artistic-infovis.pdf.
[8] For fruther discussion, see Lev Manovich, "Introduction," in Manual Lima, *Visual Complexity* (Princeton Architectural Press, 2011).

collection of important projects that visualize complex networks curated by designer and writer Manual Lima since 2004. Browsing this collection of over 700 visualizations can create an impression of almost infinite visual diversity. However, filtering them by "method" shows that many of them are variations of the same small number of visualization methods.[9] (In other words, the visual diversity of visualization field today is partly an artifact of the use of software that allows rendering the same fundamental layouts in multitude of ways.)

The endless surface variations of the small number of fundamental visualization techniques and layouts may also hide another important constant, which did not change since Charles de Fourcroy's proportional squares graphs (1782) and William Playfair's line graph and bar chart (1786).[10] Almost all information visualization techniques use a small vocabulary of discrete abstract elements: rectangles, circles, strait and curved lines, and a few others. Typically, a restricted set of a few distinct colors is used to color these elements. In other words, the visual language of graphs and visualization is the same as that of modernist geometric abstraction (1912-) and modern graphic design (1919-). Can we say that the graphs which start to first appear in the second part of the 18th century and become commonplace in scientific publications in the first part of the 19th century anticipate the development of abstract visual language in art and design a hundred years later? This is just one of many intriguing questions which waiting to be investigated by the future historians of visualization.[11]

How can we use visualization in humanities and media studies? The common sequence of steps in creating a visualization involves getting the data, organizing it in the appropriate format, and transforming it into images or animations using already existing or newly proposed technique - with the help of existing or newly developed custom software. If we want to visualize the existing data *about* cultural artifacts - for example, the lists of most popular books on amazon.com, the numbers of artworks created in different historical periods in different genres in museum collections, or the dates and locations of tens of thousands of letters exchanged by Enlightment thinkers in the 18th century (as in *Mapping the Republic of Letters* project at Stanford University)[12] - we can follow the same sequence of steps. In this workflow, the information about the lives or properties of media artifacts ends up as the familiar graphical elements of information visualization (points, lines and other graphical elements).

But can visualization also support - and hopefully augment - the key methodology of humanities: systematic and detailed examination of *cultural artifacts themselves*, as opposed to only the

---

[9] http://www.visualcomplexity.com/vc/, filter by "method."

[10] *Milestones in the History of Thematic Cartography, 1970s,* http://www.datavis.ca/milestones/index.php?group=1700s.

[11] Such an analysis will have to take into account the popularity of isotypes developed by Otto Neurath in 1920s who - while also using modernist aesthetics of simplicity and restricted geometry - also believed that isotypes will be more effective because of their iconicity.

[12] Daniel Chang, Yuankai Ge, Shiwei Song, Nicole Coleman, Jon Christensen, and Jeffrey Heer, "Visualizing the Republic of Letters", http://www.stanford.edu/group/toolingup/rplviz/papers/Vis_RofL_2009, 2009.

data about the social and economic lives of these artifacts? For example, *Mapping the Republic of Letters* projects successfully uses visualization to examine patterns in correspondence between European Enlightenment thinkers. But can visualization show all the Enlightenment letters directly rather than only dates, authors and places information – in such a way that we can both read any parts of these letters and at the same time see large-scale patterns? Or, to take another example, can visualization take further the André Malraux' idea of "museum without walls" (comparing themes and formal elements in all photographed works of art[13]) which he proposed in the middle of the 20[th] century - to allow us compare millions of professional artworks available on museum web sites, or billions of user-generated artworks on social media sites? In other words, how do we combine microscopic and telescopic vision, close reading and distant reading – "reading" the actual artifacts and "reading" larger patterns abstracted from very large sets of these artifacts?

Media vs. Data

Normally, a visualization designer works for a client who provides her with the data; the designer's job is to figure out the best way to display this data so the relationships and patterns in it become visible. However, if you a media or humanities scholar, there is no given "data" to start with. Instead, we have concrete artifacts which can come from a variety of different cultural fields: user-generated digital content, interactive design, web design, computer games, web sites, blogs, books, photographs, visual art, films, cartoons, motion graphics, graphic design, industrial design, fashion, space design, etc. This means that the default assumption of visualization that we can start with some already existing data can't be taken for granted.

There are a number of important conceptual issues involved in doing the translation from artifacts to data – here I will describe just three of them.[14]

1) The steps for translating cultural artifacts into "data" which captures their content, form, and use (reading, sharing, remixing, etc.) are not standardized - in many cases, they have to be invented and theorized. For example, what it is the "data" in the case of a web page? To make web search work, Google algorithms extract over 250 details from every web page they can find: all text, all links, fonts and colors of every paragraph, layout, etc. (The concrete details of this process are kept secret.) Would such representation of a web page be appropriate for me if I want to study and visualize the evolution of the web design since 1996, using a sample of 150

---

[13] Linda Nochlin, "Museum without Walls," *New York Times*, May 1, 2005, http://query.nytimes.com/gst/fullpage.html?res=9B0CE2DC1431F932A35756C0A9639C8B63.
[14] My discission only touches on the dimensions of this problem which I see as most relevant to visualizing media. For the theoretical and historical analysis of data practices in the sciences, see Geoffrey C. Bowker, *Memory Practices in the Sciences* (The MIT Press, 2006). For the analysis of the impact of big data on scholarly research and communication, see Christine L. Borgman, *Scholarship in the Digital Age: Information, Infrastructure, and the Internet* (The MIT Press, 2007).

billion historical snapshots of web pages from archive.org?[15] The seemingly logical answer is that this depends on the questions one I want to ask (for Google, the goal is to determine most relevant pages to the user query). However, in the case of media research, starting with well-formulated questions does not use what visualization is best at: exploring a large data set without preconceived ideas to discover "what it is there" and to find novel patterns, as opposed to only test already formulated ideas. (We can call this *exploratory visualization*.)

Even in the case of the most familiar "old media" artifacts such as printed books, it is not immediately obvious what is their "data." While typical text analysis looks at dematerialized "text" disregarding the particular formats in which it was presented to the readers, the Google search example suggests that if we are interested in reception of literature as a print medium, we do need to take into account all the details of its appearance and materiality (fonts, colors, line spacing, layout, margins, and even weight of a book).

2) Being able to translate media artifacts into data often requires specialized technical knowledge besides the domain knowledge: image processing in the case of images, computational linguistics in the case of text, audio signal processing in the case of music. To take an a concrete example from our lab, we downloaded tens of thousands of pages of *Science* and *Popular Science* magazines published between 1870 and 1922 from Google Books, and found that different sets of pages have different contrast levels. Let's say we decide that we will normalize the contrast (the decision which itself needs to be theoretically motivated). There is no single right way of doing it. There are various image processing algorithms that can be used, and each will produce a different kind of "data" as a result.

3) Translating collections of artifacts into data and then visualizing this data may "through the baby away with the water." That is, examining information visualizations of data representing aspects of cultural artifacts can lead to new understanding but it does not substitute getting insights via viewing the artifacts themselves. The last consideration is particularly important for the future of visualization in humanities, since perhaps the most important question, which is still unresolved, is how to combine distant and close readings. While all normal visualization techniques involve some reduction in order to reveal patterns, the price of this reduction is not just visibility (of new patterns) but also *opacity*, as the media artifacts with all their aesthetic richness and detail are substituted by abstract points, rectangles, lines, and curves.

In other words, the "close reading" and "distant reading" (as supported by information visualization) lead to different knowledge. Examining a text cloud visualization which shows most frequently used works in a text in order of their frequency is not the same as carefully reading the text itself. To take an opposite example, looking at the images of 9000 pages from *Popular Science* magazine (see fig. 2 below) is not the same as examining the graph that only shows metadata about these covers over time.

This discussion should explain why I gave the field of artistic visualization a prominent place in sketching the recent history of visualization. The "artistic" dimension of artistic visualization

---

[15] http://archive.org/web/web.php.

relevant for humanities is problematizing the standard visualization process, and specifically the translation of some "reality" into data. If representational cultural artifacts involve a translation – from a story, a visible world, memory, or some other type of "reality" to the signs in the artifact – visualization requires the secondary translation that maps the materiality of the artifacts into something that can be put into a spreadsheet or a database. In other words, it is a representation of a representation, a map of a map. Like any new map, it selects and omits, reveals some things and makes invisible others.

Data vs. Metadata

Having understood some of the conceptual and practical challenges of translating media artifacts into data, let us now assume that this step has been accomplished, so we can move forward in our discussion. Usually we can also assume that the media artifacts come with some metadata recorded by institutions, individuals, or software systems. For instance, over one million digital images of art, architecture and photography available via artstor.org collection are annotated with the name of the artist, year and country of creation, original size, etc. The metadata for every video on YouTube included category, tags, upload date, number of views, numbers of likes and dislikes, and so on. This metadata also needs to be problematized – rather than taking for granted the categories that it uses, we need to ask if they are meaningful. (For example, in the case of social media, is common to have metadata that specifies the countries where the users live. But what does it mean that a particular user leaves in country X? Does she leave in a capital or in a small city; was she born there and only moved there recently for school? In other words, the automatic assumption that a set of random people who happened to list the country X will have something in common is ungrounded.)

Metadata is the data about the data. Normally we assume that our goal is to study media, and the role of the metadata is to support this. However, as the amounts of media bring generated by billions of consumer devices keep growing (think of all the hours of video uploaded to YouTube every minute), the direct study of media data becomes impossible with the current methods. Instead, researchers study the metadata – because its much smaller in size than the data, because it contains structured categorical information which is easy to graph and analyze, and also because it can reveal information which can't be found in the data itself. For example, *Mapping the Republic of Letters* projects uses visualization to examine patterns in correspondence between European Enlightment thinkers. This is a typical example of the analysis where metadata itself becomes the primary object of study.

Another example comes from a project in my lab to explore images from deviantArt, the largest social network for user-generated art. We started by downloading a sample of one million images, and we also obtained metadata for these images – user screen names, upload dates, and the categories in which its creator placed each image. Having this metadata allows us to visualize the images in different ways. For instance, we can compare images submitted by different users, look at the patterns in deviantArt growth since 2000 using upload dates, and also compare images in different categories.

However, the category structure of DeviantArt is so interesting that we can study it as the artifact in its own write. Consisting from close to 2000 separate labels organized into a hierarchical tree with as many as seven levels (i.e., Customization/Skins & Themes/Linux and Unix Utilities/Desktop Environments/KDE/Styles/, Photography/People & Portraits/Spontaneous Portraits), this system presents us with a fascinating portrait of contemporary cultural imaginary. Comparing this system with the one used by museums and academics to describe visual media reveals the massive gap between the institutions of high culture and the real world. While such categories as sculpture, painting, drawing or experimental film are also present in DeviantArt, most of its categories do not have high culture equivalents (for examples, Stock images, Street Art/Stickers, or Digital Art/Pixel Art/Characters/Isometric), and yet they constitute the larger part of "non-professional art" today.

[INSERT FIGURE 2 HERE]
[Visualization of categories in dA]

In summary: What is a metadata from one perspective is the data from another perspective.

Visualizing Metadata vs. Visualizing Media

Metadata usually consists from text and numbers. We have access to a multitude of visualization techniques developed over last 300 years to represent these data types. These techniques are available in visualization, graphing and data analysis software, both free and commercial. This is another reason why practically all visualizations of humanities artifacts show only the metadata, but not the data itself.

In 2000s a few projects by digital media artists and visualization designers showed that it is possible to construct visualizations which show not only information about the images or video collections, but the images themselves. The technique used in these projects was to sample a feature film, and then display the sampled frames in a rectangular grid in the sequence corresponding to their order in the film. We created free software tools that implement these techniques, adds various options, and also makes it applicable to large collections of still images. We successfully applied the techniques introduced in these projects to variety of media forms including magazine pages, newspaper pages, comic and manga books, films, animation, and motion graphics.

Present images in a collection in a grid organized by the existing metadata such as creation or upload dates is conceptually the simplest way of visualizing an image collection.[16] We call this

---

[16] In computer science a number of researchers published papers which present more complex techniques for visualizing media colllections. However, implementing and using any of these techniques requires substantial technical knoweldge which the users in humanities and media studies do not have currently do not have. Therefore, we focused on first implementing and popularizing the techniques which are both very simple to use and very simple to explain – such

technique *collection montage*. This technique can be seen as an extension of the most basic intellectual operations of humanities – *comparison* between a small number of artifacts (typically just two) However, if 20th century technologies only allowed for a comparison between a small number of artifacts at the same time – for example, the standard lecturing method in art history was to use two slide projectors to show and discuss two images side by side – we can now compare multitudes of images by displaying them simultaneously on a computer screen. The computer graphics capacities of the current off-the-shelf computer devices (including smart phones, tablets, laptops, desktops) also allow us interact in real time with such visualizations if they show a few thousands of images only – zooming, and sorting images in different ways using of any of available metadata. But we can also construct and display static visualizations that can contain much larger numbers of images (for instance, using the software I wrote I rendered visualization which shows one million manga pages.)

This quantitative extension leads to a qualitative change in the kinds of observations that can be made. Being able to display thousands of images simultaneously allows us to see gradual subtle historical changes over tens of thousands of images, find which images are typical and which are unique, understand the patterns of similarity and difference between multiple sets of images of any size, and do many other kinds of analysis.

[INSERT FIGURE 3 HERE]
[Popular Science montage]

(Note that a rectangular grid is not the only way to display a collection. Fig, 4 below) shows another display technique that we use equally frequently. Here images are sorted in two dimensions according to their visual characteristics. This technique typically produces a view of an image collection that looks like a cloud, with image density varied in different parts of the visualization. Image with similar characteristics form tight clusters, while images with unique characteristics lie outside these clusters. The technique extends the familiar scatter plot by adding images on top of the data points. In general, we refer to the displays which show the actual images in a collection as *media visualizations* – to contrast this method with i*nformation visualization* which can only show information about the collection.)

Although collection montage is conceptually the simplest technique, it is quite challenging to characterize it theoretically in terms of where it fits into existing media forms. Given that information visualization normally starts with the text, numbers, network connection or other

---

as an image plot. For examples of research in this area, see G. P. Nguyen, M. Worring, "Interactive access to large image collections using similarity-based visualization," *Journal of Visual Languages and Computing*, v.19 n.2 (April 2008), pp. 203-224; Jing Yang, "Semantic image browser: Bridging information visualization with automated intelligent image analysis," *Proc. of 2006 IEEE Symposium on Visual Analytics Science and Technology*; Schaefer, Gerald. "Interactive Navigation of Image Collections." *FGIT 2011: Future Generation Information Technology: Third International Conference* (Springer, 2012); Gerald Schaefer, "Image browsers — Effective and efficient tools for managing large image collections," *2011 International Conference on Multimedia Computing and Systems (ICMCS.*

data type that is not "visual media" and then represents this data visual domain, is it appropriate to consider image montage as a visualization method? In this case, we start with visual domain and we end up in the same domain - starting with individual images and zooming out to see all of them. In other words, if standard information visualization translates *data into pictures*, here we translate *pictures into pictures*.

I think that calling this technique "visualization" is justified, if instead of focusing on the transformation operation of visualization (from non-visual to visual), we focus on its other key operation: layout, i.e., arranging the elements of visualization in such a way that allows the user to notice the patterns which are hard to observe in raw data. From this perspective, image montage is a visualization method. For example, the current interface of Google Books does not allow viewing thousands of pages of a magazine such as *Popular science* in a single screen, so it is hard to observe the historical patterns. However, when gather all these pages and arrange them in a particular layout (making their size the same and displaying them in a rectangular grid) using the key principle of information visualization - making everything the same on all visual dimensions except the ones where the brain will be making comparisons - these patterns become easy to see.

In its simplest form, image montage shows all images in a collection. However, with video this does not work – typically the changes between each subsequent frames are very so tiny, and showing every single frames obscures larger patterns of temporal change in content and visual form. Instead, it is more useful to sample the video, and only show the sampled frames (as this was done in the pioneering artistic visualization projects which I referred to above.) We can also apply this method to any sequential media such as newspaper pages or comic book pages. For instance, animated visualization created by my undergraduate student Cyrus Kiani uses 5930 front pages from *The Hawaiian Star* covering 1893-1912 period.[17] The animation of 5930 front pages of the newspaper published during these 20 years for the first time make visible how visual design of modern print media changes over time, in search of the form appropriate to the new conditions of reception and new rhythm of modern life. (Some of the important relevant cultural developments during this period include development of abstract art which leads to modern graphic design, the introduction of image oriented magazines such as Vogue, the spread of the new medium of cinema, invention of phototelegraph, and the first telefax machine to scan any two-dimensional image.)

The sampling procedure should not be thought about as a simple mechanical step (i.e., sample a video at 1 frame per second) or a necessary step when the data is big (as it was understood in 19[th] and 20[th] century statistics). While we can easily create a visualization showing 160,000+ frames making up a typical feature film (90 minutes = 5400 seconds = 162000 frames, assuming 30 fps rate), doing this is just not useful, as I just explained. Instead, we can think of *sampling as a creative strategy* that can be applied to any dimension of the media data. For example, in the case of an image collection, we can sample both in time (selecting every Nth image) and in space (selecting only part of every image).

---

[17] http://lab.softwarestudies.com/2012/03/visualizing-newspapers-history-hawaiian.html.

By experimenting with different ways of arranging these media samples, novel patterns can be discovered. For instance, I made a visualization which compared fist and last frames of every shot in 1928 film *Eleventh Year* by Russian directory Dziga Vertov. "Vertov" is a neologism invented by the film director who adapted it as his last name early in his career. It comes from the Russian verb *vertet*, which means "to rotate." "Vertov" may refer to the basic motion involved in filming in the 1920s – rotating the handle of a camera – and also the dynamism of film language developed by Vertov who, along with a number of other Russian and European filmmakers, designers and photographs working in that decade, wanted to "defamiliarize" familiar reality by using dynamic diagonal compositions and shooting from unusual points of view. However, my visualization suggests a very different picture of Vertov. Almost every shot of *The Eleventh Year* starts and ends with practically the same composition and subject. In other words, the shots are largely static.

I refer to the visualization method that this visualization illustrates as *remapping*. Why? Any representation can be understood as a result of a mapping operation. I am using the term "mapping" here not in a sense of production of a map of a territory but in its more abstract mathematical sense - a function that creates a correspondence between the elements in two domains. A familiar example of such mapping is projection systems used to create two-dimensional images of three-dimensional scenes such as isometric projection and perspective projection. We can also think of well-know triad of signs defined by Charles Pierce (icon, index, symbol) as different types of mapping between an object and its representation.[18]

Modern industrial media - photography, film, audio and video recoding - led to an emergence of a popular artistic strategy of using an already existing media work and creating a new meaning or aesthetic effect by sampling and re-arranging parts of this work. This strategy has been central to modern art since the second part of the 1950s. Its different manifestations include pop art, remix, appropriation art, and a significant part of media art - from Bruce Conner's very first compilation film *A Movie* (1958) to Douglas Gordon's *24 Hour Psycho* (1993), Joachim Sauter and Dirk Joachim' *The Invisible Shapes of Things Past* (1995), Jennifer and Kevin McCoy' *Every Shot / Every Episode* (2001), and numerous others.

Because many of these media art projects derive their meaning and aesthetic effect from systematically re-arranging the samples of original media in a new configuration, we think it is logical to refer to them not simply as mapping, but rather as remapping. If the original media object - a TV show, a feature film, a newspaper page, etc. - was an original media map of "reality", the art project that re-arranges its elements is a re-mapping.

---

[18] Twentieth century cultural theory often stressed that cultural repesentations are always partial maps since they can only show some aspects of the objects. However, given the dozens of recently developed methods for capturing data about physical objects and the ability to process massive amounts of data to extract new information - something which, for instance Google does a few times a day than it analyzes over a trillion web links - this assumption needs to be re-thought.

Our use of sampling and rearranging of the samples in new layouts can be conceptually related to this history. Reversely, many of the art projects that use the strategy of sampling and remapping can be retroactively understood as "media visualization." They examine ideological patterns in mass media, experiment with new ways of navigating and interacting with media, and defamiliarize our perceptions.

Although on the first glance the purpose of media visualization is simply "revealing patterns in the data," it is certainly possible to defend the position that such visualizations are more close to media art. Any remapping is a reinterpretation of the original media map, which not just teases out but also creates new interpretation and meanings.

Media visualization represents one answer to the fundamental question of how to bring together close and distant reading. Step away, and you can see larger patterns across a whole media collection. Step closer, and you can study the details of individual images.

From a semiotic perspective, media visualization breaks away from the traditional semiotics of information visualizations. The abstract elements of information visualization are symbols – signs that signify by convention. (In this, infovis can be contrasted with maps that signify by resemblance, and thus semiotically are icons.) Media visualizations show us the objects themselves, so there is no semiotic translation taking place. Rather than being symbolic representations of the objects, or their iconic maps, they are the instruments for understanding – a new epistemological technology enabled by software.

Media visualization relies on our skill to instantly to see patterns in a single image. It constructs a new image out of all images (or their samples) in a collection, arranging them in such a way that the patterns across these images can be seen as easily. Note that this method would not work with sound or text collections, since listening and reading unfolds in time. So for example, while we can arrange thousands of letters in a single high-resolution visualization as we do with images, it would not work as visualization. But arrange hundreds of thousands of images together (sorted by metadata or visual features, as described below), and the patterns are easy to see.

Adding New Metadata vs. Extracting Features

The two fundamental methods described above – using information visualization techniques to reveal patterns in the metadata, and using techniques drawn from media and digital art to display directly large media collections or their samples (and using metadata to organize the layouts) – differ in regards to what is being visualized. Information visualization shows the metadata about the media. Media visualization shows the actual data. One thing that they do share is that both methods did require adding any new information – they use already existing metadata and the contents of a collection.

I will now present two other methods that do rely on augmenting media data with new information. Both require additional work of adding this information but they differ in how this information is created.

One method that is used both in social media networks and in academic media studies and humanities is to manually add tags, or other kinds of annotations (for instance, categorical information such as in deviantArt network) using a natural language in which a researcher works in (i.e. English, Mandarin, etc.) For example, people routinely add tags to images they upload to Flickr (I am using Flickr as an example here because it popularized tagging which since then became the default feature of all social media platforms.) If Flickr's tag system employs "open vocabulary" model where any user can introduce new tags, academics usually follow "closed vocabulary" model where researchers agree on the set of tags beforehand, and then annotate a media collection using only these tags.

Many social media sites such as deviantArt also use hierarchical categorical systems to organize the media submissions (see fig. 2). These categorical metadata systems are more useful than simple tags but they also need to be approached with caution. For example, our initial investigation of appr. 280,000 image sample from the two top categories "Traditional Art" and "Digital Art" showed that while in general most users place their submissions in the appropriate categories, many do not (for example, the category "paintings" also contain many drawings.) Thus, rather than automatically assuming that categories metadata divides the data in the correct and "natural" manner, we need to think of data and metadata as two related but in the end independent entities. This view has two consequences. On the hand, the metadata itself needs to be approached as separate data set that needs to be investigated in its own right. On the other hand, automatic analysis of the data (to be discussed below) is likely to reveal clusters and groupings that do not correspond to metadata divisions.

Modern social scientists and qualitative marketing researchers use yet another way of describing a set of objects – rating objects using quantitative or qualitative scales. We can also use this approach to describe media artifacts: for instance, describing whether each image in a collection is abstract or representational on a scale of 1 to 5.

However, whether we add tags, construct our own categories and place data there, or using rating scales, all these techniques for adding new information to a media collection manually has two crucial limitations. The first limitation is that they don't work well with really large data sets. While annotating every shot of a feature film can be done in one day by a single person, imagine annotating seven billion photographs uploaded by Facebook users every month (as of early 2012) will be a real challenge if with Amazon Mechanical Turk crowdsourcing. (In the industry, some companies are able to successfully annotate large media data sets by dedicating a large stuff. For example, Pandora music recommendation engine relies on a team which rates each new song using 400 different attributes; after 10 years in business, it had a database of

800,000 songs.[19] A recent newcomer recommendation engine for art Art.sy is rumored to employ a large stuff of recent art and art history graduates who describe each artwork using a set of 800 attributes.[20])

The second limitation is using one semiotic system (natural languages) to describe another (visual media). Developing much later than senses, language complements what they do very well (capturing analog signals, differentiating between fine gradations in these signals). It allows thinking about particular and general, describing temporal relations, forming abstract categories, and differentiating between *qualities* – but it does not try to compete with the senses that are so good in capturing *quantitative* distinctions. Therefore, words that exist in natural languages to describe media aesthetics are quite limited. They can't describe the full range of variations color, texture, composition, rhythm, movement and all other analog dimensions of media. This has particular consequences for research into aesthetics of visual media, which today more than ever relies on the distinctions on these dimensions. (After abstract visual language is formulated in 1910s art, it is adopted in more and more domains –graphic design, industrial design, and architecture in 1920s, and later fashion, motion graphics, web design and UI design.)

The scale limitations mean that the manual annotation method would not work for researching the aesthetics of user-generated media if we don't want to limit ourselves to very small samples but consider patterns across large data sets (dA network contain the "modest" number 150+ million images). It can, however, work with small collections of art from the past (for instance, BBC *Your Paintings* digital archive of 200,000 paintings in UK museums[21]). However, the second limitation is always present, regardless of the size of a collection.

Instead of manual annotations, we can use well-established computer techniques to automatically process and extract information about images and video. These techniques are used in the fields of image processing, and computer vision, and many research areas such as content-based image search, video summarization, video fingerprinting, and others. Some of these techniques are known to media users – for example, face detection in iPhoto and Facebook, or smile detection used digital cameras. Other techniques remain invisible, but they form the foundation of digital media culture, as they are built in all digital media devices and applications. For example, when you take a picture with a digital camera using automatic setting, the software in the camera chip first analyzes light information captured by the image sensor, measuring gray and color values of every pixels, and then algorithmically adjusts these values to produce the image with the best contrast.

The difference between these two fundamental methods of augmenting data is not just a matter of procedure: creating new information manually or via computer image processing. The two also represent two different ways of understanding media. When we tag or annotate, its logical

---

[19] Eric Shonfeld, "With 80 Million Users, Pandora Files To Go Public". *TechCrunch*, 2/11/2011, http://techcrunch.com/2011/02/11/pandora-files-to-go-public/.

[20] "About," *art.sys*, http://art.sy/about.

[21] *BBC Your paintings*, http://www.bbc.co.uk/arts/yourpaintings.

to describe this process as *adding* additional information to the media. We can also say that we are adding new metadata to already existing metadata.

In Computer Science, the process of automatically analyzing images and video is called *feature extraction*.[22] The assumption is that computer automatically and objectively extracts the information that is already present in the images or video. The features are the statistics summarizing different types of information that can be calculated from all the pixels making up an image. Examples include average brightness, saturation and hue, number of edges and their orientations, the positions of corners, and hundreds of others. In the case of video, in addition to analyzing visual properties of every frame, temporal features such as the positions of cuts and other types of transitions between shots are also extracted (this process is called *cuts detection*).

In practical applications such as *content-based image search* (searching images by their content which in this context means both the objects in images and their visual elements such as dominant colors), hundreds of features are extracted to provide a comprehensive and yet compact representation of every image. Note that while it if well known that the choice of features has crucial effect on the success of a particular application, there is no general theory that would specify which features are to be used in different cases, so the choice of features depends on the experience of the researchers.

The two approaches – manual annotation and automatic analysis to extract features - have complementary strengths. While computers can capture the fine details of visual form, it is very difficult for them to understand the representational content of media (what images represent) - but humans can do it easily. Given an arbitrary image, we immediately detect any objects in it that have recognizable names (face, sky, house, car, etc.)

In their turn, natural languages can't capture the small differences on the visual non-narrative dimensions. For instance, try to describe using words movement patterns in tens of thousands of motion graphics works on behance.com, or other design portfolio sites.[23] While our brain can certainly compute such fine differences - over wise they would not be used universally in visual art and media - the results of these computations that drive our aesthetic and emotional responses to visual media are not accessible to the language system.

Instead of small numbers of linguistic categories, computers describe the details of visual form using real numbers. For example, let's say that we want to measure average brightness of an image. In consumer digital media brightness values are typically represented using the 256 value scale (i.e. one byte). Every pixel in an image has a gray scale value between 0 (pure black) and 255 (pure white). To measure average brightness, we add the gray scale values of

---

[22] See, for example, Mark Nixon and Alberto Aguado, *Feature Extraction & Image Processing for Computer Vision*, 3rd edition, (Academic Press, 2012).
[23] Motion graphics projects gallery on behance.net, http://www.behance.net/?field=63.

every pixel and divide them by the total number of pixels. The result is a real number (i.e. 129.54, or 178,51, etc.). Which means that our measurement scale is infinite. But even if we round off these numbers, we will still have s scale of 256 distinct values describing average brightness – which obviously provides with a much more nuanced system than the few terms available in English language (dark, medium, light).

In the same way, we can use numerical scale to characterize orientations of all lines in an image, its most prominent colors, the size and positions of all distinct shapes, and hundreds of other characteristics.

Since the two methods (manual annotation and feature extraction) complement one another, we can combine them in studying massive media data sets. For example, in our deviantArt analysis project, we run image processing software on the whole set of one million images, extracting various features from every image. We also selected a small sample of a few hundred images and tagged it manually, describing characteristics of images that computers can't capture.

Media visualizations using extracted features

While the features extracted from media collections can be explored using standard information visualization techniques such as histograms and scatter plots, they can be also used together with *media visualization* method. For example, we can sort all images according to a particular visual feature, and then render a collection montage visualization using the sorted sequence. We can also create a two dimensional scatter plot by mapping individual features to horizontal and vertical axis, and then render the images on top of the points. We find this type of visualization to be particularly useful, and we call it *image plot.*[24]

Image plots allow us to compare different image collections (or subsets of a single collection) along various visual dimensions. As an example, figure 4 shows an image plot that compares equal size samples of images from Traditional and Digital Art categories in our deviantArt sample.

[INSERT FIGURE 4 HERE]
[Imageplot comparing Traditional and Digital Art categories in deviantArt sample]

Our method of extracting visual features and then using them for media visualization draws upon existing practices in computer science – but there is one crucial difference. In computer science applications of image processing such as computer vision, content-based image search and image classification, single extracted visual features are never used by themselves.

---

[24] While the similar technique has been previosly described in a number of computer science publications, it has not beem implemented in any free or commercial software. Therefore, we developed a free software tool ImagePlot. It allows rendering of high resolution visualizations which can show very large image collections. The tool and documentation are available from http://lab.softwarestudies.com/p/imageplot.html.

Instead, hundreds of features are combined together in the hope of creating a unique "signature" for every image. If, for instance, a user wants to find all images a database similar to a particular image, the computer compares the signatures of the input image to the signatures of all other in the database images, and returns the images that have most similar signatures. (Google "search similar images" features introduced in 2009 is implemented in the similar way.[25])

This approach can be also used for media research - for instance, to identify all faces in a museum collection. However, only using computers to analyze media according to our a priori linguistic categories which can only label types of content ("people," "faces," etc.) does not use its other powerful capacity – exploring big data to see what is there, and having this exploration problematize our default understanding and assumptions. And if we want to start with a free exploration of a collection to see all kinds of patterns it can contain, or comparison between its parts, a much simpler technique is sufficient. While many of the features that can be extracted from images are not meaningful to a human observer (for instance, gray scale differences between neighborhood pixels used to characterize texture), some of them do have direct perceptual meaning. The examples of such features are contrast, the most frequently used colors, or average brightness and average saturation used in visualization in Fig. 5.

This simple but powerful technique combined with an image plot technique is our answer for how to explore image collections. The images in a collection are sorted along the dimensions defined by perceptually meaningful features. The method is simple enough so it can be taught in a single session, and it has been successfully used by my undergraduate students in a number of classes. (Obviously, we can also use existing or newly added semantic metadata in combination with image plot – as, for example, in fig. 5 which compares two subsets of our deviantArt collection using existing category information.)

It is certainly also possible to combine single visual features to arrive at more "high-level" dimensions of visual form – for instance, "calm/dynamic," or "flat/three-dimensional." However, doing this is not trivial.[26] It is not apriori clear what features best characterize such high-level dimensions, or what is the right way to combine them. In computer vision and related fields, researchers use the term "semantic" gap to describe the distance which needs to be overcomes between what computer can see - features extracted from pixel values - and the content and meaning an of image as perceived by a human. More recently, scientists introduced a related term "emotional gap" defined as ""the lack of coincidence between the measurable signal properties, commonly referred to as features, and the expected affective state in which the user

---

[25] http://www.google.com/insidesearch/features/images/searchbyimage.html; http://support.google.com/images/bin/answer.py?hl=en&p=searchbyimagepage&answer=1325808.

[26] For an example of such research, see the the following paper which investigates how low-level features can be used to describe empotional content of images: Jana Machajdik, Allan Hanbury, "Affective image classification using features inspired by psychology and art theory." MM '10 Proceedings of the international conference on Multimedia (ACM, 2010), pp. 83-92.

is brought by perceiving the signal."[27] Similarly, we can talk about "media aesthetics gap" – the distance between such low-level features and human judgments of visual form in media artifacts.

Museum Without Walls, Art History Without Names

In this chapter I looked at some of the key concepts and operations involved in the use of visualization for media analysis. These concepts are artifact, data, metadata, feature, mapping, and remapping. These concepts are basic building blocks that can be combined to form the methods that can all take us from the artifacts to their visualizations - but in different ways and with different outputs supporting different types of questions.

Once media artifacts are translated into digital data, we can decide what will be visualized. Traditional information visualization techniques are useful for exploring patterns in metadata that comes with these artifacts, new metadata manually added by researchers, or the features automatically extracted from the data representations. Media visualizations techniques originally pioneered by media and digital artists and further developed in our lab allow us to explore the patterns in images and video data itself by displaying whole collections sorted in a variety of ways. These techniques offer one solution to the fundamental question of digital humanities – how to brings together macro and micro, distant reading and close reading.

While natural languages are powerful tools for describing representational and narrative content of media, they do not work as well to describe visual form. In contrast, computers can use large numerical scales to capture nuances of form in a much more precise way. Combined with the massive media data sets now available (both digitized visual media created before 21st century, and born-digital contemporary media created by both professionals and non-professional users), this opens the door to the amazing research possibilities. Rather than only relying on small samples as media researchers did in the 20th century, we can now map histories of media aesthetics and also explore the patterns in contemporary media production, sharing and remix by analyzing billions of artifacts.

Following up on his idea of an imaginary "museum without walls" made possible by photographic reproductions of artworks, André Malraux's included 638 photographs of artworks in his book *Voices of Silence* which appeared in English translation in 1953. This was certainly a pioneering work for its time. Using media visualization, such a sample today can be expanded many times, with the numbers of images only limited by what has been digitized and what has been made available by the museums and other collections – or what can be scraped from the web. (To assemble our collection of almost 6000 images of Impressionist works that represents approximately half of the estimated number of paintings and pastels created by these artists, we scraped a number of different web sites and combined the results. For our manga project, we scraped over one million manga pages from the most popular fan manga web site together with

[27] A. Hanjalic, "Extracting moods from pictures and sounds: Towards truly personalized TV," *IEEE Signal Processing Magazine* 23(2), 90–100 (2006).

fan assigned categories.) By interactively sorting the images using both existing metadata and extracted features, displaying them in different layouts, and overlaying other historical information, we can explore their relations in ways which go beyond simple side by side comparison of a 20th century slide lecture.

This basic technique of 20th century art history was introduced by the art history Heinrich Wölfflin (1864-1945) after he became Art History Chair at Basel in 1897. He developed a teaching method of using two projectors positioned side by side in art history lectures to allow simultaneous display and comparisons of pairs of images. But this is not the only relevance of Wölfflin for our discussion. The introduction to his classical 1915 book *Kunstgeschichtliche Grundbegriffe* ("Principles of Art History") was called "Art History Names." This title reflects the ambition of art history founders - Wölfflin, Riegl, Panofsky - to analyze broad patterns of historical changes in visual representation and form on the scale of thousands of years manifested in all of the artifacts which were produced, without limiting these investigations to small sets of only important "art" objects. In *Principles of Art History*, Wölfflin writes:

> As every history of vision must lead beyond mere art, it goes without saying that such national differences of the eye are more than a mere question of taste; conditioned and conditioning, they contain the bases of the whole world picture of people. That is why the history of art as the doctrine of the modes of vision can claim to be, not only a mere super in the company of historical disciplines, but as necessary as sight itself (1932 [1915], p. 237).[28]

The broad "history of vision" advocated by Wölfflin and his contemporaries is certainly an inspiration for the use of computational analysis and visualization together with massive media collections. However, its crucial to keep in mind that this generation of researchers was limited not only by their samples and techniques of comparison, but also by the intellectual paradigms which made them read cultural artifacts as expressions of the unique characteristics describing "spirit," "mentalities," and "world picture" of different "nations."

Today, a different "art history without names" became possible – think of many millions of user-generated media artifacts and the opportunity they offer for the study of contemporary human imaginations, including both their "content" and the patterns of imitation, diffusion and innovation on a global scale. Media visualization methods allow us to explore such massive collections without a priori reducing them to small number of categories as Wölfflin and others had to do. And rather than assuming that media created by users which have similar demographic profiles has something in common (to translate Wölfflin's assumptions in contemporary terms), we can instead use the combination of feature extraction and media visualization to find clusters of similar media objects, and then see if they correspond to user demographics or any other existing categories.

---

[28] Quoted in Michael Hatt, Charlotte Klonk**,** *Art History: A Critical Introduction to Its Methods* (Mancester University Press, 1998), p. 66**.**

Ultimately, visualization can help us to question our existing metadata labels and ways of dividing the objects of study, showing that that every narrative and map we construct is only one possibility - as Bruno Latour puts this, "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."[29]

Acknowledgments

References

"About." *Art.sy*. Web. <http://art.sy/about>.

Archive.org. *Archive.org*. <http://archive.org/web/web.php>.

BBC Your paintings. Web. <http://www.bbc.co.uk/arts/yourpaintings>.

Berry, David, ed. *Understanding Digital Humanities*. New York: Palgrave Macmillan, 2012. Print.

Borgman, Christine. *Scholarship in the Digital Age: Information, Infrastructure, and the Internet*. The MIT Press, 2007. Print.

Bowker, Geoffrey. *Memory Practices in the Sciences*. The MIT Press, 2006. Print.
Burdick, Anne, Johanna Drucker, Peter Lunenfeld, Todd Presner, Jeffrey Schnapp. *Digital Humanities*. Cambridge, Mass.: The MIT Press, 2012.

Chang, Daniel, Yuankai Ge, Shiwei Song, Nicole Coleman, Jon Christensen, and Jeffrey Heer. "Visualizing the Republic of Letters." 2009. Web. <http://www.stanford.edu/group/toolingup/rplviz/papers/Vis_RofL_2009>.

---

[29] Bruno Latour, *Tarde's Idea of Quantification,* in Mattei Candea, ed., *The Social After Gabriel Tarde: Debates and Assessments* (Routledge, 2009).

Friendly, Michael and Daniel J. Denis. *Milestones in the History of Thematic Cartography, Statistical Graphics, and Data Visualization*. Web. <http://datavis.ca/milestones/; http://www.datavis.ca/gallery/index.php>.

Gold, Matthew, ed. *Debates in the Digital Humanities*. Mineapolis: University Of Minnesota Press, 2012. Print.

Hanjalic, A. "Extracting moods from pictures and sounds: Towards truly personalized TV." *IEEE Signal Processing Magazine* 23(2): 90–100. 2006. Print.

Hatt, Michael and Charlotte Klonk**.** *Art History: A Critical Introduction to Its Methods*. Mancester University Press, 1998. Print**.**

Hayles, Katherine. *How We Think: Digital Media and Contemporary Technogenesis*. Chicago: University of Chicago Press, 2012. Print.

Nochlin, Linda. "Museum without Walls." *New York Times*, May 1, 2005. Print. http://query.nytimes.com/gst/fullpage.html?res=9B0CE2DC1431F932A35756C0A9639C8B63.

Latour, Bruno. "Tarde's Idea of Quantification*." The Social After Gabriel Tarde: Debates and Assessments*. Ed Mattei Candea. Routledge, 2010. Print.

Lima, Manual. *Visual Complexity*. Princeton Architectural Press, 2011. Print.

Lima, Manual. *Visualcomplexity*. Web. <http://www.visualcomplexity.com/vc/>.

Machajdik, Jana and Allan Hanbury. "Affective image classification using features inspired by psychology and art theory." Proceedings of the 2010 international conference on Multimedia. ACM, 2010. Pp. 83-92. Print.

Manovich, Lev. "What is visualization?" *Visual Studies*, 26.1 (2011): 36-49. Print.

Manovich, Lev. "Cultural Analytics: Visualizing Cultural Patterns in the Era of 'More Media'." Milan: *Domus*, 2009.

Manovich, Lev. "Data Visualization as New Abstraction and Anti-Sublime." *SMAC!,* 3 (2002). Print. <http://lab.softwarestudies.com/2008/09/cultural-analytics.html>.

Manovich, Lev. "How to compare one million images?" *Understanding Digital Humanities*. Ed. David Berry. New York: Palgrave Macmillan, 2012. Print.

Manovich, Lev. "Introduction." *Visual Complexity.* By Manual Lima. Princeton Architectural Press, 2011. Print.

Manovich, Lev. "Media Visualization: Visual Techniques for Exploring Large Media Collections." *Media Studies Futures*. Ed. Kelly Gates. Blackwell, 2012. Print.

Moere, Andrew Vande. "About the Information Aesthetics Weblog." *Infosthetics.com*. Dec. 2004. Web. <http://infosthetics.com/information_aesthetics_about.html>.

Motion graphics projects gallery. *Behance.net*. Web. <http://www.behance.net/?field=63>.

Nguyen, G. P., M. Worring, "Interactive access to large image collections using similarity-based visualization." *Journal of Visual Languages and Computing*. 19. 2 (April 2008): 203-224. Print. <http://staff.science.uva.nl/~giangnp/Pubs/Pdf/2006/NguyenWorringJVLC06.pdf>.

Nixon, Mark and Alberto Aguado. *Feature Extraction & Image Processing for Computer Vision*. 3rd edition. Academic Press, 2012.

Ramsay, Stephen. *Reading Machines: Toward an Algorithmic Criticism*. University of Illinois Press, 2011. Print.

Schaefer, Gerald. "Interactive Navigation of Image Collections." *FGIT 2011: Future Generation Information Technology: Third International Conference*. Springer, 2012. Print.

Shonfeld, Eric. "With 80 Million Users, Pandora Files To Go Public." *TechCrunch*, February 11, 2011. Web. <http://techcrunch.com/2011/02/11/pandora-files-to-go-public/>.

Software Studies Initiative. ImagePlot software. Web. <http://lab.softwarestudies.com/p/imageplot.html>.

Software Studies Initiative. *Software Studies Initiative*. Web. <http://www.softwarestudies.com>.

Viégas, Fernanda and Martin Wattenberg. "Artistic Data Visualization: Beyond Visual Analytics." *Proceedings of the 2nd International Conference on Online Communities and Social Computing*. Springer-Verlag Berlin, 2007. Print. <http://www.research.ibm.com/visual/papers/artistic-infovis.pdf>.

Yang, Jing, Jianping Fan, Daniel Hubball, Yuli Gao, Hangzai Luo, William Ribarsky. "Semantic image browser: Bridging information visualization with automated intelligent image analysis." *Proc. of 2006 IEEE Symposium on Visual Analytics Science and Technology*. 2006. <http://viscenter.uncc.edu/sites/viscenter.uncc.edu/files/CVC-UNCC-06-02.pdf>.

Yau, Nathan. *Visualize This: The FlowingData Guide to Design, Visualization, and Statistics.* Wiley, 2011. Print.