Interactive urban analytics: time, place & data

Balmer, Jeffrey [1] Dou, Wenwen [2] Ribarsky, William [2]
Sauda, Eric [3]

Keywords: urbanism; visualization; representation; interactivity; digital humanities

Abstract The emerging field of visual analytics offers powerful new tools with which to understand complex built environments. In this paper, we demonstrate how to apply these interactive techniques to both temporal and spatial data in cities, using Rome, Italy, with its many interconnected histories, as a case study.

The technologies comprising temporal visual analytics create an *intertwined* history. Such a history begins with an overarching temporal narrative of the entire historical record of important events. We define 'event' as a meaningful occurrence at a particular *place* and *time*. An event in weather history could be a great storm, while one in religious history could be the building of a church. An intertwined history places events from various histories into the same time structure. We have chosen Rome because of its long history, abundant scholarship and overlapping geographies. Our methods include:

- Identification of events and sub-events and arrangement into hierarchical tree-structures.
- Correlation with GIS-based spatial structure.
- Derivation of meaningful topics from collection of written documents using the *Stanford Modeling Toolbox*.
- Generation of time based 'scaffoldings' as event outlines.
- Collection of imagery based on maps, engravings, photographs, and automatic extraction of architectural imagery from perspectives.
- Use of *Theme River* and *Event Identification* visualizations.
- Design of interactive multi-view visualizations.

Visual analytic interfaces such as the one that we demonstrate have several important features that advance their usefulness.

1. Data is displayed in multiple windows using different forms of analysis. These windows are linked in such a way that data highlighted in one representation is simultaneously shown in multiple iterations.
2. Changes in the viewpoint in any window can allow a user to explore the data interactively. This means that a user can explore the data, coupling the computational capacities of the computer with the human ability to perceive patterns and meaning.
3. Analytic techniques allow the presentation of potential clusters of information to the user. Rather than operating upon a static archive, these techniques allow for possible conjectures of knowledge embedded within the huge silos of data.

1. School of Architecture, University of North Carolina, Charlotte; USA
2. Department of Computer Science, University of North Carolina, Charlotte; USA
3. Digital Arts Center, School of Architecture, University of North Carolina, Charlotte; USA

Introduction

Throughout the humanities, emerging digital technologies have begun to offer researchers powerful new tools with which to assemble and analyze archival materials. In this paper, we trace the development of increasingly sophisticated and interactive techniques within the nascent field of *digital humanities*. We present a visual analytics system designed to assist scholars in the study of urban centers, and select Rome, Italy, with its many interconnected histories, as a case-study.

Digital Humanities is a loosely defined term that refers to any intersection of the humanities and computation. Many projects identified with digital humanities involve the generation of digital archives and associated data and search structures. While these efforts are useful in providing easier access to information, they are not fundamentally different from their analog counterparts.

Visual Analytics is a recently defined field within computer science. It takes as its fundamental premise that analysis is better accomplished as a symbiosis of the computational power of computers with the pattern and sense-making capacity of human users. For the purposes of this paper, we will focus on projects that engage two issues; the theoretical structure by which a body of work is conceptualized and the interaction of users with the computer.

Precedents

Franco Moretti's *Graphs, Maps, and Trees* (Morelli, 2007) is an early attempt to introduce ideas of quantitative models to the study of literature. Moretti rejects the study of exceptional literary occurrences and focuses attention on literary production as a whole. He uses three simple mathematical constructs to organize this research. In one study, he maps the total number of novels over time using simple line graphs over time in several countries. Subsequently, he uses diagrams to examine the dynamics of a novel, and trees as a structure to show the evolution of the form and techniques within a novel. Each of these operations is intended as a strategy to destabilize conventional readings of 'great' literature. The statistical objectivity of his models are meant to provoke tension with more orthodox forms of literary studies

Lev Manovich's *Style Space* (Manovich, n.d.) studies artistic production, employing simple scatter-plot diagrams to plot their distribution. He begins with two sample sets of Van Gogh paintings, one executed in Paris and the other in Arles. Using software to obtain a median value for saturation and brightness in each painting, he then plots the results against the two variables. These results can then be interpreted for patterns of grouping, evenness of distribution and cross comparison of sets of images. Other graphs analyze reverse skew value, plotting painting values against standard art historical categories, projecting trend lines against sets of data points and graphing the number of elements in each painting over time. Manovich makes no claim about the comprehensiveness of such analyses, instead seeing them as a tool to reveal otherwise hidden patterns, and allowing researchers to posit meaningful patterns.

These approaches reveal several shortcomings with existing approaches: There is little or no coordination among the various graphs; we are forced to infer the relationships between them. More seriously, there is no opportunity for user interaction with the data. Both of these features are fundamental to the visual analytic approach that we propose below. To facilitate the analysis of enlarged surveys of data, the necessity of time-based analysis, and a common geospatial frame of reference, we propose to use a specific set of techniques from visual analytics as a tool allow for searching and formulating questions within the humanities.

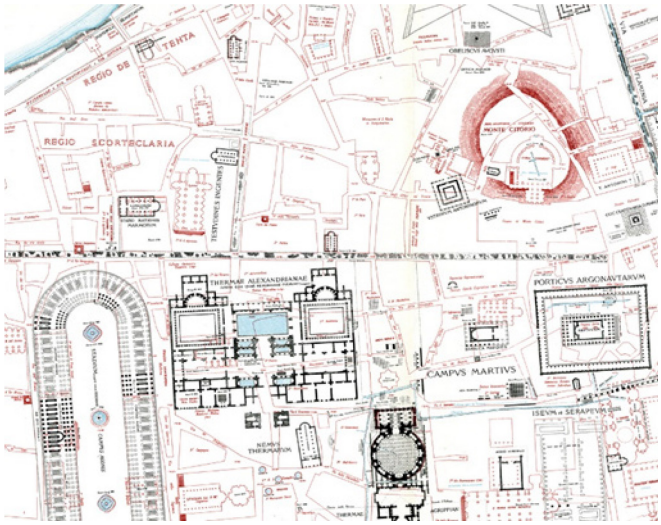

Fig. 1 Lanciani map of Rome, 1901, among the first attempts to create a spatio-temporal database

Visual Analytics: Case-study

Let us imagine how visual analytics might assist the work of an architectural historian. Our historian is undertaking research on the interior iconography of Sant'Andrea al Quirinale, a church begun in 1658 for the Jesuit seminary on the Quirinal Hill in Rome, and designed by Gian Lorenzo Bernini. Because of the significance afforded this church by generations of scholars, our historian will find a rich array of secondary sources with which to begin his or her understanding of its significance: as a space of worship, as evidence of the priorities of its patrons, and as the work of one of the most celebrated and influential artists of the Italian Baroque.

Should our historian wish to pursue archival research, with an attempt to yield new insights from primary sources, they would find similarly ample materials at hand. Assuming that this research aims to establish some form of context for the pictorial, sculptural and architectural program of the interior of the church, a historian could begin their work with a number of approaches, depending on their prior interests and expertise. They might begin by examining Bernini's entire oeuvre, attempting to place Sant'Andrea in the larger arc of his entire artistic production. Similarly, research might focus within the context of churches constructed at the same time and place, say Roman churches begun in the 1650s, or perhaps those constructed in the Quirinal district during the 17th century. Another approach would be to examine churches dedicated to Saint Andrew, or those commissioned by the Jesuit Order, within reasonably narrow parameters of time and place, to allow for an exhaustive examination of all relevant archival records.

Traditionally, as the frame of reference grows wider, so too does the scope of work entailed in an exhaustive examination of primary sources. Exhaustive coverage of a subject is demanded of scholarship, such that it may be deemed both relevant and authoritative. Seldom has such work been carried out to the satisfaction of a historian's peers. Lewis Mumford's sweeping chronicle of Western civilization is one example (Mumford, 1968). But such herculean efforts will remain the province of those rare individuals with either super-human stamina or a rare access to the support (financial and temporal) required of such work. Even then, such works are surveys that adopt a measure of superficiality as a necessary premise. The pitfalls of such attempts are well documented: A historian may establish conclusions based on their necessary reliance on secondary sources, which runs the risk of being biased, flawed, or merely incomplete.

Visual Analytics: Capacities

A visual analytic approach could track not just the history of a single building, but of entire districts in Rome, while preserving a researcher's direct access to primary sources. The Quirinal Hill, largely depopulated after the fall of the Empire, was resettled following the completion of the Acqua Felice in 1586 by Sixtus V. With the re-establishment of a ready supply of water to this elevated quarter of the city, inhabitation was once again possible. The impetus for construction was driven by the expansion of institutions affiliated with the Church in the wake of the Counter Reformation. As such, detailed records of the transfer of ownership, of construction, of patronage contracts for the outfitting of churches and seminaries, and the activities that these buildings sponsored, are available in both Church and state archives in Rome[1]. Such data, assembled within a map-based matrix would reveal a fine-grained spatio-temporal record of patterns over centuries.

Further, such findings could be supplemented by text-based sources. Such data is already capable of being analyzed, not just as traditional notions of concordances or keywords, but with the context and syntax of phrases and entire sentences retained. Advances in analyzing vast tracts of text allow richer and more instructive patterns to emerge. What patterns might these systems support, and what conclusions might they enable? In our focus on the church of Sant'Andrea, several scenarios arise. Texts referencing the church, whether scholarly, bureaucratic, or merely anecdotal, might present patterns of interest to a scholar with the use of various filters. Might a time-line based analysis show spikes in usage among one or more types of text sources? Would the originating locale of the text vary over time? Patterns of secondary concordances (the pairing of the church with mention of other artifacts, ideas, or individuals) could yield both expected and unanticipated results.

Visual Analytics: Techniques

Scholarship centered on a city like Rome has focused on a range of architectural, political, social, economic, demographic and geographical dimensions. Such studies typically adopt one particular lens; an architectural historian may examine a building in detail, but will have less complete and secondary knowledge of geographical, economic and social factors. While some disciplines are inherently focused upon historical issues, there is always a tension between syntactic and diacritic approaches, and no field has a set of conceptual tools to deal with multiple histories from disparate disciplines.

The technologies comprising a temporal visual analytics approach to research would underwrite what we call an *intertwined* history of Rome. Such a history would begin with an overarching temporal narrative of the entire historical digital corpus of important events. Here we define 'event' as a meaningful occurrence at a particular place and time. A meaningful event in weather history could be a great storm, while one in religious history could be the building of a church. An intertwined history would place events from various histories into the same time structure.

Events can be arranged and displayed in multiple ways simultaneously. For example, using satellite and Doppler radar observations, (Yue Jiang et al., 2001) a complex history of a large-scale hurricane can be automatically arrayed as a series of sub-events (formation, approaching land, storm surge, severe wind and rainfall, etc.) that provide a compact signature for the whole story of the hurricane. (Li Yu et al., 2010) Alternatively, interactive event timelines can be displayed as shown in Fig. 2. Here, news events have been extracted automatically from CNN over an entire month. Each event-bubble shows all the news stories for a particular main event, with the thickness of the bubble proportional to the number of related and concurrent news stories. Selecting a bubble at a particular point produces a video of the main news story.

[1] In addition to the Vatican Library, and the Jesuit archive, substantial records for property transactions, construction documents and property assessment records are contained in archives and libraries administered by both the city of Rome and the Italian Ministry of Culture.

Once we have the collection of events and sub-events over time, they may also be arranged into hierarchical tree structures. The next step is to combine such tree structures with the spatial structure of GIS. The incorporation is first spatial. A standard GIS Spatial Analysis in Macro-ecology is used to divide and subdivide the surface of the virtual world into spatial cells, and the locations of the events, city boundaries, buildings, social activities, etc. (for an urban history) are placed in their respective cells. To this combination, we add the time hierarchy over the events for this spatial region. The appropriate level for Rome might be a spatial cell of, say, 20 KM on a side, centered on the Forum. The organization of the time structure over this spatial cell is in terms of events. All major events occurring in this cell are inserted at their appropriate places along the time dimension.

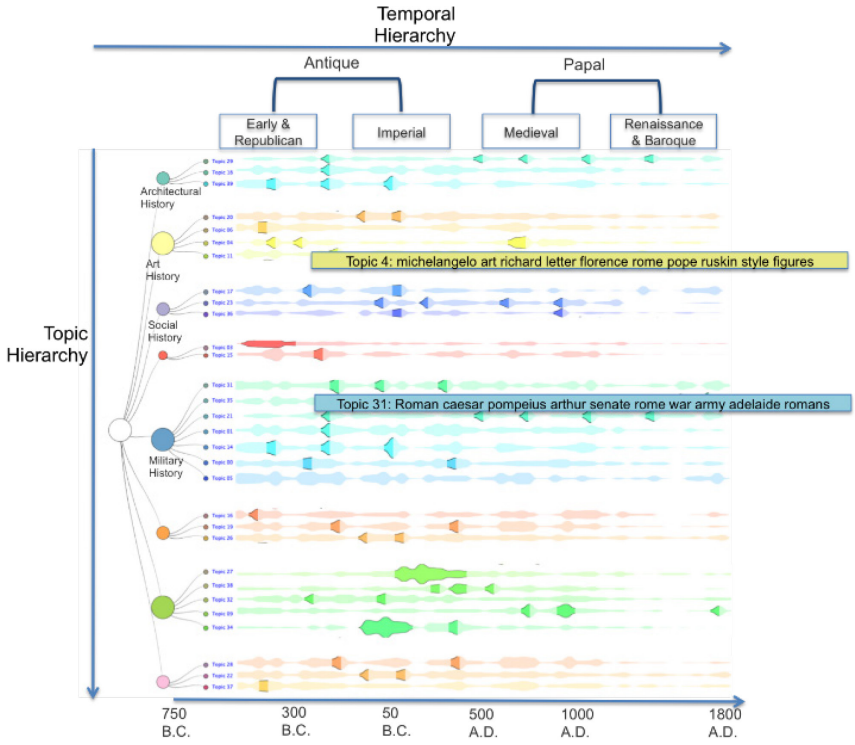

Fig. 2 Interface Mockup. The topics and events are organized hierarchically. Grouping occurs wherein one branch contains topics containing similar attributes. Topic hierarchies can be interactively modified to adapt to users' conceptual models of topic spaces. Events (clusters) are then identified. Here, the x-axis represents a timeline of Rome. Topic content (shown in text boxes) is developed via user interpretation

Recently developed methods can successfully derive meaningful topics from collections of written documents. In its simplest iteration, these automated, data-driven methods analyze documents as a so-called 'bag of words', without regard to linguistic relationships.

Although detailed contextual meanings may be lost, the methods are quite flexible and can be used to derive topics from the underlying content that, under human analysis, are meaningful. Such methods allow the collection of documents to provide patterns of content to an analyst, including how the individual documents cluster into related groups and the nature of their relationships. In a recent NY Times article, Robert Nelson, director of the Digital Scholarship Lab at the University of Richmond, described how data-mining can be successfully applied to topics in the complete collection of stories reported by a Southern newspaper during the Civil War, revealing trends over time, and in response to key events that might remain otherwise obscured. (Nelson, 2011)

Fig. 3 Illustration of geospatial view. Coordinated with the topic view shown in Fig. 2, the x-axis denotes time. The dots in the map denote events extracted from the textual sources that happened in the area. The size of the dots indicates the magnitude of the event, i.e. how many digital records documented a particular event. The geospatial view provides an overview of events in both space and time. Coordinated with the topic view, one can derive a comprehensive understanding of an event based on what, when, where, and whom (from entity extraction)

One method we have used for text-based analysis (Dou et al., 2011) is based on a Bayesian approach (Blei et al., 2003) that automatically learns a set of topics for a document collection determined by a set of keywords, and then assigns appropriate topics to each document in the collection. The Stanford Modeling Toolbox[2] makes such tools accessible. Such approaches can be applied to full-text documents (papers, essays, reports, books--including works of fiction), although books will need to be broken into chapters, Fig. 4 shows the topics derived from a collection of NSF proposals using this approach. (Dou et al., 2011) In this case, twenty topics were automatically extracted from the proposal collection, and displayed as vertical lines in the top part of the figure. Each individual proposal is a horizontal line graph plotted showing its proportion for each topic. The topics, some of which are labeled, well describe the themes of the NSF directorate examined (IIS in this case). The bottom of the figure is derived from ThemeRiver timeline software, where each colored band shows the ebb and flow of particular topics across a horizontal timeline.

Coupled with these tools, for an intertwined visual analytics of Rome, we must add means to extract entities indicating time, location, actors (people or groups), and references (e.g., building names, neighborhood names, etc.) for each document section related to given topics. Many techniques already exist for this sort of geo-temporal analysis. Among the most powerful of these tools are those built around FactXtractor, (Pan et al., 2007) developed at Penn State. A collection of interactions permit one to filter by topic, time period, related proposal, among others.

How might we convert disparate collections of events into intertwined histories? We begin with information already organized and available. Compact historical syntheses, such as Richardson's *New Topographical Dictionary of Ancient Rome* (Richardson, 1992), or Talbot's

[2] Stanford Topic Modeling Toolbox. http://nlp.stanford.edu/software/tmt/tmt-0.3/

Barrington Atlas of the Greek and Roman World. (Talbot, 2000) By digitizing these texts, their relevant content could be analyzed with topic modeling and entity extraction for additional topics, related events, and for relations among topics. The initial, high-level event extraction could be achieved manually or semi-manually in support of significant meaning; main events would reveal themselves readily. Even if incomplete in terms of main events, this initial timeline would soon be filled in by the topic modeling and event extraction step upon analysis of subsequent texts.

Additional texts containing syntheses of social history, ecclesiastical history, architectural history, etc., would then be added, building up a series of 'scaffoldings' to form broad outlines of events for each of these histories. With this initial structure, we can then cast our net wider, to other books and documents, using the initial collection as a 'training set' for the broader sets of texts, to enrich existing topics and events, and to uncover new ones. Ultimately, we would consider this set of scaffoldings together, seeing how they intertwine and connect with one another. Simply seeing events from different histories in close proximities of time and place may be sufficient to infer relationships, including those of possible cause and effect. Once identified, such events may then be explored more deeply. Ideally, topic modeling could be tuned to automatically reveal potential relationships in more detail, and relate events from different histories into a comprehensive history (Fig. 2, Fig. 3).

The relative significance of events among these histories are assessed by examining text collections, which can identify events mentioned across multiple collections. We have already developed a process similar to the one proposed here, although in a much more limited form, in a project studying the impact of NSF programs and the proposals they fund. (Dou et al., 2011) Here we needed to follow programs that covered particular research areas over time (e.g., databases or visualization in computer science), even as the names and to some extent the focus of the programs change across time. We overcame this problem by tracking the underlying scientific principles expressed in research papers published, where fields such as databases or visualization have permanence over time and evolve slowly. For each field (e.g. visualization), we determined primary journals and conferences, applying topic modeling to the papers (over time) from this collection. Our analysis revealed main topics and trends over time. Topics that appear across several subjects are cross-disciplinary, as are proposals that have multiple salient topics. This approach can also infer the impact of a particular program on the underlying science. (We can see how a research area evolves after a particular funding program is begun.) It can also reveal what research areas a program is actually funding, as distinct from what the program *says* it is funding.

Vast collections of imagery associated with a city like Rome are widely available. Visual data from the past century is primarily photographic imagery and digital media. From the period we refer to as Papal Rome, this imagery is primarily drawings, paintings, and prints. From the ancient city, extensive archeological evidence survives. Recent advances in what visual analytics call 'image semantics' have led to breakthroughs in automated analysis of image content. These methods have already been applied to digital image collections such as Flickr.

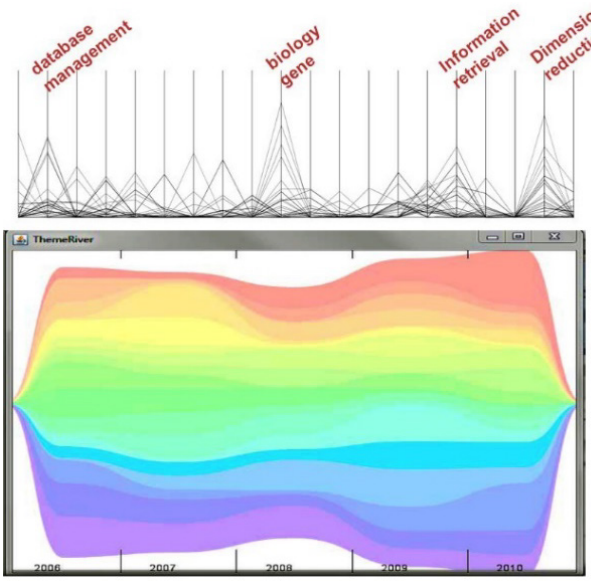

Fig. 4 Topics from text collection and ThemeRiver timeline

Digital Historiography

Visual analytic interfaces such as the one that we propose have several important features that advance their usefulness to researchers. First, data is displayed in multiple windows using different forms of analysis. These windows are linked in such a way that data highlighted in one representation is simultaneously shown in multiple iterations. This range of views present different aspects of the data, and allow a researcher to visualize information in several forms side by side.

Second, changes in the viewpoint in any window can allow a user to explore the data interactively. This means that a user can explore the data and test what meaning he or she may see within. This interactivity uses the computational capacities of the computer in tandem with the human ability to perceive patterns and meaning. This complementary relationship uses the best abilities of both machine and man. Finally, the specific time-based techniques that we propose are uniquely suited to the study of history, and allow any user to visualize and explore complex spatial, social, and material relationships over time. Interconnections among events in widely dispersed records can be juxtaposed and assessed for possible significance.

Third, analytic techniques allow the presentation of potential clusters of information to the user. Rather than operating upon a static archive, these techniques allow for possible conjectures of knowledge embedded within the huge silos of data. These are presented to the user not as answers or fixed meanings, but rather as a set of possibilities. Our approach always assumes that the user will be uniquely suited to determine the meaning and significance of these patterns, while the machine is uniquely suited to uncovering possible connections within huge databases. A further extension of this approach is the ability to capture the insight of an expert and make it available as a possible technique to other researchers.

The emerging capacities of visual analytics offer an entirely new set of tools to assist our view of history, one freed from canonical reading and allowing for research as a standard mode of

exploration, in what Joshua Sternfeld[3] has called digital historiography. Or, as Paul Arthur has said, "*Records are no longer fixed, but dynamic. The record is no longer a passive object, a 'record' of evidence, but an active agent playing an on-going role in lives of individuals, organizations and society*". (Arthur, 2008)

Finally, we acknowledge the unease with which many in the humanities may view the encroachment of the computer. If our aim were to use digital technology to supplant human judgment, this suspicion would be well founded. What we propose is not an abdication of human thought and judgment, but rather an extension of its reach and power. Our approach recognizes the unique abilities that humans can use in conjunction with the unique ability of computers to deal with huge and complex challenges.

[3] Sternfeld, Joshua, 'Thinking Archivally: Search and Metadata as Building Block for a New Digital Historiography'

References

Moretti, F. (2007). *Graphs, Maps, Trees: Abstract Models for Literary History*, New York: Verso.

Manovich, L. (n.d.). *Style Space*, http://lab.softwarestudies.com/2011/08/style-space-how-to-compare-image-sets.html

Mumford, L. (1968). *The City in History: Its Origins, Its Transformations and Its Prospects*. New York: Mariner.

Yue Jiang, T. et al. (2001). Acquisition and Display of Real-Time Atmospheric Data on Terrain. EG/IEEE EuroVis 2001. New York: Springer-Verlag.

Yu, L., Lu, A., Ribarsky W., and Chen, W. (2010). Digital Storytelling: Automatic Animation for Time-Varying Data Visualization. *Computer Graphics Forum*, Vo. 29 (7), 2271-2280.

Nelson, R.K. (2011). Of Monsters, Men—And Topic Modeling. *New York Times* (May 29, 2011).

Dou, W., Chang, R., Wang, X., and Ribarsky, W. (2011). ParallelTopics: A Probabilistic Approach to Exploring Document Collections. Accepted for publication, *IEEE* VAST 2011.

Blei, D., Ng, A., Jordan, M. (2003). Latent Dirichlet allocation. Journal of *Machine Learning Research*, 3:993–1022.

Dou, W., Chang, R., Xiaoyu, W., Ribarsky, W. (2011). ParallelTopics: A Probabilistic Approach to Exploring Document Collections. Accepted for publication, IEEE VAST 2011.

Pan, C.C. et al. (2007). TexPlorer: An application supporting text analysis, *IEEE* VAST, 205-6.

Richardson, L., (1992). *A New Topographical Dictionary of Ancient Rome*, Johns Hopkins UP, Baltimore

Talbot, R., (2000). *Barrington Atlas of the Greek and Roman World*. Princeton: Princeton University Press.

Arthur, P. (2008). 'Exhibiting History: The Digital Future'. reCollections: The National Museum *of Australia.*