The Whole Story: Building the Complete History of a Place

William Ribarsky
The Charlotte
Visualization Center
UNC Charlotte
ribarsky@uncc.edu

Eric Sauda
School of
Architecture
UNC Charlotte
ejsauda@uncc.edu

Zachary Wartell
The Charlotte
Visualization Center
UNC Charlotte
zwartell@uncc.edu

Jeffery Balmer
School of
Architecture
UNC Charlotte
jdbalmer@uncc.edu

Abstract

In this paper, we discuss how one can build a visual analytics system to comprehensively describe a place throughout its many interconnected histories. We discuss the needed 4D data structure, the analytics techniques, and the interactive visualizations. This combination of automated and interactive techniques can be brought together into a new, powerful capability. We focus on the example of Rome and, more specifically, on its architectural/cultural history.

1. Introduction

The founding of cities coincided with the rise of human civilization and systems of law, written language, organized religion and economic exchange. History was always written with awareness of a particular place, whether it was its political history, military history, or a description of its great men. Great cities (Babylon, Athens, Rome, Alexandria) evolved into city-states, regions around the city, and eventually into empires. Only since the renaissance have nation states arisen and histories been written about the politics, governments, and military outcomes of these larger and more diffuse entities. Lately, there have arisen many other threads of the history of a place: social history, cultural history, history of common life, architectural history, history in terms of natural events that affect the lives of the population, and so on. It is obvious that these histories are all intertwined and interconnected. If we could read them together, we would gain extraordinary new insights and knowledge about how and why a place changed.

In this paper, we investigate how this whole story of a place can be built by considering a concrete example, the city of Rome. Rome has been alive for nearly three millennia and has been a prodigious center of political, social, and religious life for over two millennia. Think of all that has been written about Rome, all the paintings and drawings that have been created, all the maps and views of the city that have been made over the centuries. Think of all the records kept by imperial, state, and local governments and by the church over the millennia. Think of the

archaeological evidence that is still being uncovered, that exists in layers under present churches and public buildings. Think of the geographic and geologic record of the Tiber as it ebbed and flowed over the centuries. Think of the story of the surrounding lands and all the weather that has enveloped the city and its effects on human records and geology. Almost all of these things have a location, a region, or a range in space and time. They can thus be incorporated (and organized) into a virtual world model, and then cause and effect relations inferred via their proximity in time and space and via the nature and meaning of these things.

There is an incredibly large and bewildering amount of information about Rome in words, pictures, maps, and other media. Whole libraries are devoted to just one portion of this story. How can we make the story complete; make it look like one integrated thing? The methods of visual analytics, coupling interactive visualization with data analytics and analytical reasoning, can show the way. Now available are a collection of interactive data analysis methods that fit well with interactive visualization. These methods include topic modeling for large text collections [1], image understanding for large image collections [2], extraction of 3D geometry (e.g., building location, shapes, volumes) from perspective images [3], GIS methods to extract and organize meaningful layers from maps, and other computer vision techniques that, for example, can understand unstructured video or precisely locate detailed snapshots of scenes.

This organized information, plus numerical or other data on physical events, provides the comprehensive picture of Rome we seek. We can then construct a 4D virtual world model (4D GIS) to bring everything together in its proper relationship. In this case, the result will ultimately be a comprehensive, accessible organization of very nearly everything that has been recorded and observed about Rome. This space-time GIS will have the efficiency of a normal GIS but will be much more extensive and powerful. One might ask of it, for example, what major events occurred in Trastevere in a given time period and whether they were related to events across the Tiber in the imperial district. Or the system could automatically set up an

event time line for the whole area or any subarea. But this time line would not be the usual inert representation; it could be made to show all events of a certain importance, not just, say, political or military events. For example, one could see a major flood or plague event near social or political events that might indicate upheaval. This might reveal a relation that was hidden before. The timeline would be interactive and responsive. Major events could be opened up to indicate a detailed history. Selection of any event would indicate its spatial extent and, if dynamic, its movement on the spatial view. Choosing related events at different times for the same area would show interesting patterns of development or evolution. Selection of major building events on a timeline representing an area of the imperial city could show how medieval, renaissance, and modern buildings were erected on the imperial sites, sometimes directly on top of earlier structures, or around them, as happened with some medieval churches. These structures could be shown simultaneously, using semi-transparent representations of the buildings. The result would be comprehensive, novel views of the history of an area. One could follow multiple, intertwined bundles of narrative arcs for any aspect of the city.

In this paper, we discuss how one might build this visual analytics system to comprehensively describe a place throughout its history. We will discuss the needed 4D data structure, the analytics techniques, and the interactive visualizations. We will focus on the example of Rome and on its architectural/cultural history. However, the approach proposed here could be extended to the other histories of Rome and, indeed, to other cities or nations. Our purpose in this paper is not to present a complete system; but rather to present the ideas and methods required to build such a system.

2. Relevant Work

Two key requirements for our vision are efficient storage and query of 3D spatial-temporal data and automated extraction of events that are meaningful to the end-users (historians, etc.). Both the historical data and the events themselves must be efficiently stored and queryable. A database *access method* is the set of algorithms and data structures (or the *index*) that support these queries. Database access methods support data sets that exceed in-core capacity. Many temporal access methods (TAM) and spatial access methods (SAM) are available and are now implemented in temporal [4] and spatial databases [5]. A large number of spatial-temporal access methods (STAM), which tightly integrate time and space, have been developed [6] and many domain specific STAMs have been developed in computer graphics and

imaging, robotics, CAD and GIS [7][8][9]. Database access methods (or *AMs*) by definition handle out-of-core data sets and their index resides out-of-core. However, TAMs, SAMs and STAMs in these other domains often only support in-core data sets or, more subtly, they may dynamically load the data but require that the entire index fit in-core. These restrictions are necessary to support rendering frame rates or robot planning algorithms that execute hundreds to millions of specialized queries (such as collision detection or visibility) every 16 to 100 milliseconds.

STAMs do not yet appear in commercial or complete open-source DBMS's. Further, database STAMs remain an active area of research. Examples include handling geometric objects that are highly mobile and streaming [10] or that contain dynamic topology [11]. At a recent workshop to set National Geographic Intelligence Agency research agenda, a call was made for a GIS with time structure as rich as the spatial structure [12].

In this paper, we introduce a relatively unexplored approach to time structuring in terms of events. We show that this is a general approach that is applicable to many types of events and data. We then describe how this can be merged into an urban visual analytics approach that includes topic modeling and entity extraction to extract the whole story of a place.

3. Events and 4D GIS

To tell the whole story of Rome, there is the further issue of how to automatically generate and present to the end-user an over-arching temporal narrative of the entire historical digital corpus. One way to do this is to develop a narrative structure based on important events. For this purpose, an event is defined as a meaningful occurrence at a particular place and time. Extending the idea of multiple, intertwined histories presented in the Introduction, a meaningful event in, for example, a weather history could be a great storm whereas one in religious history could be the building of a church. An intertwined history would place events from both these histories into the same time structure.

These events can be extracted more or less automatically in multiple ways. For example, recent weather events can be extracted from satellite or Doppler radar observations [13] or from simulations using 3D feature extraction techniques. Thus the whole history of a large scale hurricane can be automatically set up 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 [14]. Other events, including weather events that predate modern observational methods, will be described in texts or picture/text combinations

The techniques described in the next sections can be used to attach meaning to and make a history out of these events. Although these automatic techniques usually provide good event identification, the meanings of the events are not always clear. In these cases, a user can interactively attach a meaningful label or category to the event. Since higher level events are much more

Once we have the collection of events and sub-events over time, we can organize them into a hierarchical tree structure; see Figure 2. (Undoubtedly a comprehensive structure will be more complicated than a tree with, for example, links between events across categories. However, we will use a tree as a first step because it captures a high order organization (events, sub-events, etc.) that is meaningful to the user. Then links between main events within or between trees can be made explicitly.) The question is then how to combine this tree structure with the spatial GIS structure. The incorporation is first spatial. A standard GIS SAM is used to divide and subdivide the surface

To insure fast access to relevant information, we will maintain both the spatial and the event hierarchies. The events are compact, so this hierarchy will not require much space or overhead, while the GIS hierarchy will insure the most efficient access to the potentially much more voluminous spatial data. (The relative compactness is due to the fact that the events are a high-level abstraction of the detailed underlying data). With this structure one could filter by events and time or by spatial area. Using this structure, it would be quite easy to construct interactive time lines annotated by main events along their time range that could then be expanded to reveal sub-events.

Figure 2 shows the high-level data organization. The Historical Digital Corpus is all the text, images, video, 2D and 3D models, and GIS layers referenced by the software system. These may be stored in their own DBMS's of varying types; in plain, flat files; or in domain specific AMs (such as graphics data formats). In the figure, the Narrative Event Hierarchy is the hierarchical, narrative event structure that is visually presented to the user. The Event STAM stores Events by their time period and 3D location. The Geometry Representation STDAM (spatio-temporal-*detail* access method) provides access to the geometric representations (i.e. spatially located computer graphics models) associated with the Historical Corpus elements. Scalable graphics data structures require level-of-detail aspects [16] which the rendering algorithm uses to only process a graphics object at the resolution required by the current viewpoint. Hence, the Geometry Representation AM includes this additional level-of-detail component.

Access method design involves trade-offs among the type of queries that the method implementation will execute with optimal speed. The Geometry Representation (or *G.R.*) STDAM is optimized for performance of real-time, interactive rendering and 3D interaction techniques (selection, manipulation, collision detection, etc). At a low-level, no one STDAM is optimal for all types of rendered geometry.

Geo-located information including map icons and glyphs, 3D meshes, volumetric data, point clouds, 2D video, image-based rendering, etc. all have specific spatial, temporal and/or detail access methods [16][17]. Also, many graphics AMs do not page their index or page their data. The G.R. STDAM itself, however, will page its data and its index, but its data leaf nodes will reference specialized graphics AMs and render code modules appropriate for a given node. In all cases, the goal is to support *interactive* visualization since the data will be so rich with unexpected relationships that an exploratory end-user experience is necessary.

Figure 2. Data Structures and Access Methods

The G.R. STDAM entries are Spatio-Temporal Entries (STEs) defined by a spatial extent and a time period. An STE also contains a reference to a paged coarse geometric representation, a reference to the detailed AM index and data set, a reference to a render code module, and a reference to the set of the generating data in the Corpus.

STAMs have a broad design space with dozens of different independent dimensions [7][8][9]. One aspect is whether the temporal index is built above the spatial structure, the spatial structure is built above the temporal structure or the indices are interleaved. The Event STAM will be organized with time on top of space. The G.R. STDAM will be organized with space on top of time. However, the individual graphics AMs associated with a given Spatio-Temporal Entry will be organized in whatever way is best for the Entity's geometry representation. Events in the Narrative Event Hierarchy will be instances of a Spatio-Temporal Entity. Each Event's Spatio-Temporal Entity will be indexed by both the Event and the G.R. AM's.

Spatial subdivision in both the Event STAM and G.R. STDAM will use an octree. The Event STAM will place an interval tree [18] on top of its octree. In the G.R. STDAM each octree cell contains an interval tree. This interval tree contains a reference to every Spatio-Temporal Entry whose spatial extent intersects

the octree cell and whose spatial extent just fits into a box the size of the cell's parent cell.

The rendering traversal of the STDAM would use HLOD methods [19][20], but when reaching an STE, the STE's render module is executed and passed information such as the current time instant and the projected size in pixels of the STE's extent. For efficiency, the renderer should take advantage of frame-to-frame coherence to avoid re-traversing the entire octree and the entire temporal interval tree for the octree nodes that pass the coarse cull and LOD tests. Again, the render modules for each STE would be highly specialized to the particular geometric representation of the STE.

4. Topic Analysis and Entity Extraction

Recent methods can successfully derive meaningful topics from a collection of documents. These automated, data-driven methods analyze documents as a bag of words without regard to linguistic relationships. Although detailed higher 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. The methods can be applied to documents of any type in any field. Combined with other automated methods, they allow the collection of documents to "tell the observer what they are about" including how the individual documents cluster into related groups and the nature of their relationships. Indeed, these methods can be applied to different languages; all that is needed is an appropriate dictionary. A recent New York Times article demonstrated how this topic modeling approach can be successfully applied to the trends in topics in the complete collection of stories reported in a Southern newspaper during the Civil War, revealing trends over time and in response to key events that would be quite hard to discover otherwise [21].

An approach that we have used successfully to analyze NSF proposals over time and also do full text analysis of relevant papers [1] is based on a Bayesian approach [22] that automatically learns a set of topics for a document collection in terms of a set of key words and then assigns appropriate topics to each document in the collection. The Stanford Modeling Toolbox [23] makes such tools available. These approaches can be applied to full-text documents (papers, essays, reports, books--including works of fiction), although long documents such as books will need to be split via chapters, etc.

Figure 3 shows the topics derived from a collection of NSF proposals using this approach [1]. In this case, 20 topics were automatically extracted from the proposal collection. The topics are vertical lines in the top part of the figure; each individual proposal is a

horizontal 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 a themeriver where each colored band shows the ebb and flow of particular topics over time.

Figure 3. Topics from text collection and ThemeRiver timeline.

Coupled with this, we need the 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 exist for this sort of geo-temporal analysis. Among the most powerful of these tools are those built around FactXtractor [24], developed at Penn State. With appropriate geo-cross-referencing, these topics can be used to help form the events described in the last section. A collection of interactions permit one to filter by topic, time period, related proposal, and other ways.

But how do we convert collections of events into histories, and how do we make these histories rich and complete? We do not have to start these from scratch; there is a lot of information already organized and available. A compact historical synthesis of the best scholarship, such as Michael Grant's history of antique Rome [25], or a more recent historical atlas, will provide a rich timeline of main events and sub-events. These can be extracted and then relevant parts of the book and of other books can 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 done manually or semi-manually in support of good meaning; the main events would be straightforward to

find and might even be listed in the book (or in chapter headings). Even if incomplete in terms of "main events", this initial timeline would be filled in by the topic modeling and event extraction step upon further analysis of the historical text collection.

One can do the same thing for other histories, finding leading books or other historical overviews that contain syntheses for social history, ecclesiastical history, architectural history, and so on. The idea is to build a series of "scaffoldings" containing the events for each of these histories. Once we have this initial structure, we can cast our net wider to other books and documents, using the initial collection as a "training set" for this broader collection so that we can enrich existing topics and events and find new ones.

Ultimately, one wants to consider this set of scaffoldings together, seeing how they intertwine and affect each other. Of course, just seeing events from different histories that are close together in time and location, perhaps having correlated trends over time, may be enough to infer a relationship, including a possible cause and effect relationship. These events can then be explored more deeply. In addition, it would be very useful if the topic modeling could automatically reveal potential relationships in more detail and combine main events from different histories into a comprehensive history.

Relative importance of events across histories can be assessed by looking at their text collections to see which events are mentioned in more than one collection. Moreover, if topics or parts of topics associated with an event appear across histories, it is a further indication of the event's importance. In this way a comprehensive hierarchy of events across histories could be built.

In fact, we have already developed a process similar to the one proposed here, although in a much more limited form, in our project to study the impact of NSF programs and the proposals they fund [1]. Here we need to follow programs that cover particular research areas over time (e.g., databases or visualization in computer science) but the names and to some extent the focus of the programs change. We surmount this problem by tracking the underlying science as expressed in research papers published, because fields such as databases or visualization have permanence over time and evolve slowly. For each field (e.g. visualization) we determine the main journals and conferences and then apply topic modeling to the papers (over time) from this collection. We can thus find main topics and trends over time. Topics that appear across disciplines 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 show what research areas a program is actually funding (through topic analysis of the proposals funded and comparison to the topic analyses for the research fields), which may be somewhat different from what the program *says* it's funding. We have applied this approach to visual analytics and have found trends over time and the effects of individual funding programs.

5. Image/Map Analytics

As described further in the next section, huge amounts of imagery associated with Rome are available. More recently, this has been photographic imagery and other media, but in the antique and especially the papal periods, this imagery was drawings, paintings, and prints. A lot of work has been done on image semantics that lead to image understanding. In this approach, instead of low-level image processing, semantics of imagery are determined for higher level image understanding (e.g., whether they depict beach scenes, urban scenes, scenes with people pursuing certain activities, traffic scenes, scenes with flowers, types of animals, etc.). These techniques are robust (i.e., they can recognize beach or traffic scenes even under conditions of different lighting, viewpoint, or activity), and new topics for automated understanding can even be user-created by selecting sets of training images for the desired topics. Further, the image understanding algorithms coupled with user interaction can very efficiently remove junk images and provide high quality tagging. These methods have been applied to image collections such as Flickr.

For the intertwined histories described here, these methods would need to be extended to analysis of drawings, paintings, and prints. Some work has been done with drawings--especially map drawing understanding, which will be of substantial use here [27]. This map content can be associated with geographic views of relevant events. In addition, there are new methods that successfully consider image and text content together [28]. These methods will permit automatic derivation of much more meaningful descriptions for images (e.g., from captions or closely associated text). Accurate analysis of drawings, paintings, and prints will depend on the legibility of the original and will be affected by the artist's style. However, interactive visualization of the clustered images after automated analysis can quite efficiently clarify miss-assignments and miss-groupings.

Finally, we need methods to automatically extract architecture from perspective (or birds-eye) views of Rome and its neighborhoods painted or captured at different times in its history. The Vatican museum

alone has thousands of such images. Debevec et al [29] developed a semi-automated system for modeling and rendering architecture from photographs. Myriads of later methods were developed which further automate the creation of high detailed 3D models from both historical photos and paintings. El-Hakim et al [3] demonstrate the applications of a range of techniques for combining both photos and paintings to generate 3D models of historical structures. More recently Schindler et al. [30] combined automated 3D reconstruction with a historical probabilistic model than can ingest a set of historical photographs of a cityscape with partially known dates and then construct 3D models of buildings to determine the most plausible temporal ordering of the original photographs. With these methods and the event structuring described above, an urban visual analytics system can be built.

6. Focus on Rome

Rome has been as a city for 2,500 years, and has been the center of three distinct social orders: Antique Rome, spanning the period from the republic to the largest empire of the ancient world, Papal Rome, the seat of ecclesiastical and secular power during the renaissance and baroque eras, and Modern Capital, Rome as the center of the Italian state. It is this layering of culture, art, urban form and architecture that forms much of the character of the city.

The kinds of information that are available from each era vary with the interest in the culture of the era, how well records are preserved, and, in the case of Rome in particular, interpretations by one era toward its predecessor. All of these factors make the application of this visual analytic approach particularly fruitful for Rome, given the need to understand an almost bewildering variety of records and ideas. Indeed, much of the scholarship that centers on Rome proceeds by ignoring layered complexity and instead focusing on one particular time period or type of record. For example, experts will be able to examine the political and economic motives and consequences of papal correspondence from the 16th century, but will not have tools to understand this together with other information about Rome during the same period.

With the applications of visual analytic techniques described here, scholars who embrace the idea of "digital humanities" will be able to understand and compare information from disparate sources, potentially finding previously hidden relationships between economics and aesthetics, between transportation systems and theology, between geography and urban form. The shift in tone and structure of the 16th century papal correspondences may be linked in meaningful ways to economic, social,

geographic and technical issues that have been previously impossible to understand.

We have identified three data sources for each period that illustrate the range of data types and at the same time contrast the information that may be available from other periods. For Antique Rome, we pursue mytho-poetic history as presented by leading authors of roman history, the catalog of inscriptions on public buildings, and the archeological record of building type and locations. For Papal Rome, we pursue the bureaucracy of an ideology as preserved in the Vatican written records, a catalog of renaissance painting animated by newfound interest in an accurate representation of nature, and the building record of new projects aimed at restoring Rome as a capital. For the Modern Capital, we pursue GIS for the city of Rome and Lazio, the National Institute of Statistics for census data and Flickr images geo tagged to Rome.

6.1. Antique

Although the number of primary sources documenting the history of any city from antiquity is dwarfed by historical records of more recent periods, Rome offers us an exception for several reasons. As the seat of a military and economic entity that dominated Western civilization for over 500 years, Rome offered its historians a unique wealth of both patronage and primary sources. Even at the dawn of this Imperial period, the city had already accumulated five centuries of legislative, bureaucratic and literary documents. As a city that at its peak numbered a million inhabitants, more than 20 times the size of any Western city before it, the range and depth of recorded human activity is unparalleled in the ancient world. And as self-conscious inheritor of the classical traditions of Hellenic culture, the Imperial city fostered archives and a culture of scholarship that assembled a sweeping record of all antiquity. Sources include the histories of Livy, Plutarch, Sallust, Suetonius and Tacitus, the correspondence of Cicero, Seneca and Varro, the treatises of Pliny, Vitruvius and Frontinus, and the musings of Virgil, Ovid, and Horace. Although most of these histories are more mytho-poetic than scholarly, these texts, and those they have spawned throughout subsequent centuries, form a uniquely influential contribution to Western culture.

As a more durable corollary to the relatively ephemeral texts inked on papyrus and parchment, the extensive collection of inscriptions on buildings (Figure 4), monuments and tombs both public and private throughout Rome's Republican and Imperial periods is unequaled in the ancient world. In addition to the extant remains, subsequent catalogs from the Papal period provide us with a uniquely detailed record

of primary sources that range from the august (official dedications and bureaucratic minutiae) to the perfunctory and illicit (graffiti and political sloganeering). As with the more traditional literary forms that survive, ancient Rome provides for nearly two millennia of secondary and tertiary sources contingent upon its wealth of these primary materials.

Figure 4. Inscription, Arch of Titus. Rome.

Rome provides an incomparably rich and layered archeological record, from the early Iron Age settlements excavated atop the Palatine Hill to the late Antique remains of a metropolis in transition, there are few, if any, cities that rival Rome for both the range of archeological artifacts, and the depth and scope of scholarly discourse. Materials yielded from the archeological record of Rome are both image and text-based in nature, and have been extensively published and discussed via the broad array of scholarly journals published throughout the past one hundred years. This record includes current and ongoing activity across the historic city center and its environs.

In addition, of course, there is a wealth of post-antique texts that describe political, social, urban, military, and economic life during this period. These texts refer to and organize antique sources. Our system will ingest these texts to provide the backbone of an integrated history of this era.

6.2. Papal

As Rome transitioned from the temporal to the ecclesiastical center of the Western world, it continued its status as an administrative center of unique breadth and reach. As such, even with the fall of its imperial status, it upheld its capacity for recording and archiving a broad range of human activities, both sacred and profane, throughout some of the most turbulent events in Western history. The libraries within Rome itself contain a bureaucratic record of almost two thousand years, in its capacity as the seat of Catholicism. Extant written records are contained in the Vatican library, and historical collections elsewhere in the city (the Biblioteche Angelica, Casanatense,

Hertziana, and Vallicelliana, to name a few). In addition to the singularity of its ecclesiastical holdings, the extent of library collections of literature, philosophy, history, and the natural sciences is rivaled by few cities in the West.

Arguably the greatest patron of the arts throughout Western civilization, the Church has embellished the city of Rome with an inimitable wealth of artistic production. There are extensive collection of works that depict the city, in whole or in part, in ways that are sufficiently naturalistic to aid in the reconstruction of the city. The diversity of works and subjects, from the early drawings of Duperac and Palladio, through the paintings of Pannini and van Wittel, the etchings of Vasi and Piranesi, and the maps by Bufalini and Nolli (Figure 5) provide us a broad spectrum of pictorial sources that document the city throughout this period. Collections include the Vatican and Capitoline Museums, the Barberini, Borghese, Colonna, Corsini, Doria-Pamphili and Spada Galleries, as well as smaller collections both public and private throughout the city. Images originally made in the city are also contained in collections elsewhere in the world.

Figure 5. Letarouilly's Plan of the Campodoglio, 1840.

As with the ancient city, the built environment of the Papal city is both unique and ubiquitous throughout the historic center of Rome and its environs. As the 'headquarters' of Catholicism, the city is endowed with an unrivalled diversity of sacred architecture spanning almost two millennia, from the earliest Christian catacombs to the ongoing construction of contemporary churches designed by internationally acclaimed architects. Additionally, the city features a unique diversity of building types sponsored by the broad range of ecclesiastical, diplomatic and charitable institutions overseen directly or indirectly throughout the history of the Roman Curia. Particularly in the wake of the Counter-Reformation, religious orders (particularly the Jesuits) and related social institutions provided the city a unique mix of 'ecclesiastical infrastructure' seen nowhere else. Detailed records of

changes in ownership and occupation of buildings in the historic center of Rome in both Papal and Civic archives have the potential to yield significant analytic insights into the development and transformation of the city's inventory of historic structures.

6.3. Modern

The contemporary history of Rome is in most ways typical to that of other cities of economically developed countries. One of the most interesting aspects of studying Rome is the ability to intersect modern technologies and data sources with the past.

Like all modern cities, Rome is awash with sources of data. We will focus on three sources that typify the much larger set of data available. One underlying similarity of all three is the use of spatial location, a result of both a technological basis and the development of modern cartography.

The Sistemi Informativi Territoriali (SIT) is the name for the Italian Geographic Information System. Information for Rome and the Italian government are available at a publically accessible web site [31]. This is generally supported using ArcGIS or compatible software, with the limits on data display and interactivity inherent in that system. While huge amounts of data are available and are all geo referenced, there is no established methods of correlating data that are not commensurable. One of the goals of this project will be to find ways to develop meaningful connections with other forms of data and, more importantly, to register these data with information over the whole history of Rome.

Some work has already been attempted in Rome using GIS both for the identification of historic sites and with overlay of historical maps [32]. Our system will provide a comprehensive framework for these fragmentary attempts.

The Istituto Nazionale di Statistica (Istat) is the Italian National Statistical Institute, which since 1926 has conducted the census. The census is conducted every ten years; the last complete census was 2001. Complete statistics are available on-line [33], including demographic, economic and social information at the census block level.

The proliferation of images is a hallmark of contemporary life. The painstaking and rare paintings and frescoes of the papal era have been replaced by a flood of images. As Paul Virilio noted, as a result of "...the industrial proliferation of visual and audiovisual prostheses...we are looking at the rapid collapse of mnemonic consolidation" [34].

An extreme form of this is on-line image sites, such as Flickr, which contain billions of images. A more recent trend is using geo tags for images, which leads

to an archive that is specifically tied to Rome. A recent check shows almost 500,000 such images (Figure 6). Some tools, such as Google's Picasa, have taken initial steps toward embedding images tied to events within virtual spatial models. The interests of tourism skew this record of images, but they form a powerful resource of current and modestly historical desire.

Figure 6. Flickr with 480K geotagged photographs.

7. Putting It All Together

Once the system described here is put together, what could be achieved? Integrated time lines of main events will not be hard to construct by hand; entity extraction with at first a well-chosen set of sources can then fill in details about these events and also relationships between them. The event timeline in Figure 3 and the annotated Topic river in Figure 1 would be linked to a geographic view. As events for a topic are selected, the geographic view would reveal the location and spread for that event/topic. Conversely, a region of the city could be selected and only events/topics for that region displayed.

Imagine an historian studying the Counter Reformation. Normally, she would primarily study the chronology of events presented from one or two archival sources, keeping in mind the political, economic, architectural and geological histories. This system would allow her to explore these intertwined histories interactively, linking event histories together, seeing which events fall together in sequence, and uncovering relationships both spatially and topically.

Imagine an architect studying one particular location in a city, but being able to see not just the current form of the urban context, but also all the underlying events that co-locate on the site. This could include the accumulation of structures and neighborhoods (extracted from city views rendered at different times) that rose and fell alongside the main buildings, how the city built up in layers over time, and even how particular buildings (such as the Pantheon)

changed over time in appearance and function. The history of geological events such as flooding, the economic history of the site, the demographic history of the site can all be examined simultaneously. What will this record reveal about the uses of certain locations or neighborhoods now versus Renaissance or Imperial times, including the subtle ebb and flow of uses? There are architects who know small parts of a city in this manner, but it is an unusual and uneven knowledge. Urban visual analytics will provide access to the full wealth of information.

Imagine following the demographic history of Rome over its history, rising to a city of a million inhabitants in Imperial times and falling to a few thousand residents, living in an outdoor museum during the Middle Ages, and then recovering its grandeur and its population during the Renaissance and expanding into Modern times. Intertwined histories will reveal how these demographics affected and were affected the social and political history of Rome.

Of course, we need digitized records in order to apply our automated methods. Many contemporary records are digitized. In addition, the Vatican Library has undertaken the digitization of 80,000 of its historic manuscripts (with images and illustrations) and will place the archive on the Web. Large portions of painting collections are also digitized. Secondary and tertiary sources (e.g., texts from journals and recent books) are also digitized. A related question is whether the texts are scanned or have detailed word recognition. A lot of the available older manuscripts are probably scanned, so a character recognition step would have to be applied (with quality of recognition dependent on the quality of the original imprint). Fortunately there is a wealth of English texts in archives around Rome; Rome has been a focus of English-speaking scholars for centuries. Although our automated methods could be used on other languages with a dictionary providing keyword and caption translations, it is certainly easier and richer to have English texts. We conclude that although the record is not complete, there is plenty of digital information to find events, do event structuring and topic modeling, and analyze collections of images.

8. Conclusion and Next Steps

In conclusion, we have described a system of automated and semi-automated methods that can be brought together to create the complete story of a place. We and others have shown how particular methods can successfully produce results that fit into the overall story. Bringing them together in the integrated approach described here will undoubtedly yield discoveries and new knowledge, possible

examples of which are described in the last section. Interactive visualization plays a key role both in providing meaning to retrieved events and topics and in uncovering trends, relationships, and patterns that would be hard or impossible to see otherwise.

Our plan is to begin constructing main events and event histories for political, religious, urban (including architectural), and military Rome. We will also add historical details about main people movements, plagues, and population demographics. We will then begin generating appropriate topics for these histories.

9. References

- [1] Wenwen Dou, Remco Chang, Xiaoyu Wang, and William Ribarsky. ParallelTopics: A Probabilistic Approach to Exploring Document Collections. Accepted for publication, IEEE VAST 2011.
- [2] J. Fan, W. Ribarsky, R. Jain, "Social computing for collaborative image understanding", Invited Paper, *IEEE Multimedia*, vol. 18, 2011.
- [3] El-Hakim, S.F., J.-F.c. Lapointe, ois, and E. Whiting. Digital reconstruction and 4D presentation through time. in ACM SIGGRAPH 2008 talks. 2008. New York, NY, USA: ACM.
- [4] Snodgrass, R. T. (1999), *Developing Time-Oriented Database Applications in SQL*, Morgan Kaufmann.
- [5] Kothuri, R. K. V.; Ravada, S. & Abugov, D. (2002), Quadtree and R-tree indexes in oracle spatial: a comparison using GIS data, Proc, ACM SIGMOD 2002, New York, NY, USA, pp. 546--557.
- [6] Theodoridis, Y., et al. Specifications for efficient indexing in spatiotemporal databases. in Scientific and Statistical Database Management. Proc. Tenth International Conference on. 1998.
- [7] Rigaux, P., M. Scholl, and A. Voisard, Spatial databases with application to GIS. 2002, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- [8] Goodchild, M.F., M. Yuan, and T.J. Cova, Towards a general theory of geographic representation in GIS. *Int. J. Geogr. Inf. Sci.*, 2007. 21: p. 239-260.
- [9] Samet, H. Sorting in space: multidimensional, spatial, and metric data structures for computer graphics applications. ACM SIGGRAPH ASIA 2010 Courses. New York, NY.
- [10] Mokbel, M.F. and W.G. Aref, SOLE: scalable on-line execution of continuous queries on spatio-temporal data streams. *The VLDB Journal*, 2008. 17: p. 971-995.
- [11] Jiang, J. and M. Worboys, Event-based topology for dynamic planar areal objects. *Int. J. Geogr. Inf. Sci.*, 2009. 23: p. 33-60.
- [12] National Academy of Sciences Workshop. New Research Directions for the National Geospatial Intelligence Agency (May, 2010, Washington, DC).
- [13] Yue Jiang, T., et al. Acquisition and Display of Real-Time Atmospheric Data on Terrain. EG/IEEE EuroVis 2001: Springer-Verlag, New York.
- [14] Li Yu, Aidong Lu, William Ribarsky, and Wei Chen. Digital Storytelling: Automatic Animation for Time-Varying Data Visualization. *Computer Graphics Forum*, Vol. 29 (7) pp. 2271-2280 (2010).
- [15] D. Luo, J. Yang, M. Krstajic, J. Fan, W. Ribarsky, and D. Keim. EventRiver: Interactive visual exploration of constantly evolving text collections. *IEEE Trans. On Visualization and Computer Graphics* (2011).
- [16] Luebke, D., et al., Level of Detail for 3D Graphics. 2002, New York, NY, USA: Elsevier Science Inc.
- [17] Akenine-Möller, T., and Haines, E. *Real-time Rendering*, A.K. Peters, Natick MA, 2002.
- [18] Edelsbrunner, H., Dynamic Data Structures for Orthogonal Intersection Queries. 1980, Inst. Informationsverarb., Tech. Univ. Graz, Graz, Austria.
- [19] Erikson, C., D. Manocha, and I. William V. Baxter. HLODs for faster static and dynamic environments. in I3D '01: Proc. of the 2001 symposium on Interactive 3D graphics. New York, NY, USA: ACM.
- [20] Yoon, S.-E., B. Salomon, and R. Gayle, Quick-VDR: Out-of-Core View-Dependent Rendering of Gigantic Models. *IEEE Transactions on Visualization and Computer Graphics*, 2005. 11(4): p. 369-382.
- [21] Robert K. Nelson. Of Monsters, Men—And Topic Modeling. *New York Times* (May 29, 2011).
- [22] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [23] Stanford Topic Modeling Toolbox. <http://nlp.stanford.edu/software/tmt/tmt-0.3/>
- [24] Pan, Chi-Chun, Jaiswal, Anuj R., Luo, Junyan, Robinson, Anthony A., Mitra, Prasenjit A., MacEachren, Alan M., Turton, Ian A., TexPlover: An application supporting text analysis, IEEE VAST 2007 pp. 205-6.
- [25] Grant, Michael, History of Rome, Prentice Hall, New Jersey, 1978.
- [26] T. Shaw and P. Bajcsy, "Automated Image Processing of Historical Maps", SPIE (7869), pp. 7869-9, 2011.
- [27] H. Samet and A. Soffer, "A Legend-Driven Geographic Symbol Recognition System", ICPR94, pp. 350-55 (1994).
- [28] H. Luo, J. Fan, Y. Zhou, "Multimedia news exploration and retrieval by integrating keywords, relations and visual features", *Multimedia Tools and Applications*, vol.51, pp.625-648, 2011.
- [29] Debevec, P.E., C.J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. Proc. ACM SIGGRAPH '96. New York, NY, USA: ACM.
- [30] Schindler, G. and F. Dellaert. Probabilistic temporal inference on reconstructed 3D scenes. *IEEE CVPR*, 2010.
- [31] "Italy-Nationwide Data," Geo Community, accessed 6/9/11, <http://data.geocomm.com/catalog/IT/datalist.html>
- [32] K. Lelo and M. Travaglini, The GIS-based historical atlas of Rome, presented at the XXth International CIPA Symposium, 26 Sept.–1 Oct. 2005, Torino, Italy.
- [33] "Databases and information systems," Istat, accessed 6/9/11, http://en.istat.it/dati/db_siti/.
- [34] Virilio, Paul, *The Vision Machine*. Bloomington: Indiana University Press, 1994.