CRITSPACE: AN INTERACTIVE VISUAL INTERFACE TO DIGITAL COLLECTIONS OF CULTURAL HERITAGE MATERIAL

A Dissertation

by

MICHAEL NEAL AUDENAERT

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

December 2011

Major Subject: Computer Science

CritSpace: An Interactive Visual Interface to

Digital Collections of Cultural Heritage Material

CRITSPACE: AN INTERACTIVE VISUAL INTERFACE TO DIGITAL COLLECTIONS OF CULTURAL HERITAGE MATERIAL

A Dissertation

by

MICHAEL NEAL AUDENAERT

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Approved by:

Chair of Committee,	Richard Furuta
Committee Members,	Frank M. Shipman
	Tracy Hammond
	Maura Ives
Head of Department,	Hank Walker

December 2011

Major Subject: Computer Science

ABSTRACT

CritSpace: An Interactive Visual Interface to Digital Collections of Cultural Heritage Material. (December 2011)

Michael Neal Audenaert, B.S., M.S., Texas A&M University

Chair of Advisory Committee: Dr. Richard Furuta

Cultural heritage digital libraries have become an important and prominent tool within humanities scholarship, offering increased expressive power for representing complex networks of relationships and the ability to use computational tools and interactive environments to help researchers ask new questions. While digital libraries offer tremendous advantages for publishing the final products of scholarship, in the words of Bradley and Vetch, "as they currently are delivered, do not intersect terribly meaningfully with the process of scholarly research." In this work I investigate how scholars use visually complex source documents—materials where access to a visual representation of the original object is required and present a prototype system, CritSpace designed to facilitate scholarly engagement with digital resources. Rather than creating a one-size-fits-all application, CritSpace is a web-based framework for building interactive visual interfaces that support scholarly use of digital libraries. The theory and design behind CritSpace is based on a formative study of the work practices of scholars from different disciplines and prior research in field of spatial hypertext. To illustrate a concrete example of using CritSpace and to evaluate its usefulness, I

conclude with a case study that walks through the process of deploying CritSpace to support work in a specific scholarly domain, textual criticism and presents a summative usability study of the tool. The results of this study show that CritSpace is effective at supporting textual criticism. More significantly, they also indicate that the innovations added in CritSpace promote the intensive analysis of visual material in addition to knowledge organization and structuring.

DEDICATION

This work is dedicated to my family. First and foremost to Michelle who has patiently waited through many late nights, put up with my grumblings, and generously not held me accountable for the numerous “I’ll be done by” dates that came and went. I would not have finished this work without her support. And to my parents who have motivated and supported my quest for knowledge without ever pressuring.

ACKNOWLEDGEMENTS

I would like to thank my committee chair, Dr. Furuta, and my committee members, Dr. Shipman, Dr. Hammond, and Dr. Ives for their guidance and support throughout the course of this research.

Thanks also Dr. Warren for his generous support in organizing the final user study and for making himself and the staff of the Center for New Testament Textual Studies available for this project. I also want to extend my gratitude to the participants in both of the studies reported here. Their time and effort is much appreciated.

George Lucchese's work on the initial prototype of CritSpace as an REU student was simply amazing. He was directly responsible for much of the early implementation and design work that went into the system and for several innovations that have, unfortunately, been left to future development.

Finally, thanks to the students of the Center for the Study of Digital Libraries for the many fruitful conversations and debates about a wide range of topics over the past few years—especially to Unmil Karadkar for his mentoring on graduate student life, leadership and service at the University, and research in general.

TABLE OF CONTENTS

CHAPTER Page

LIST OF FIGURES

LIST OF TABLES

CHAPTER I

INTRODUCTION

Cultural heritage digital libraries have become an important and prominent tool within humanities scholarship. The specific advantages of digital libraries to many areas of practice within the humanities are now well-documented. These include both increased expressive power for representing complex networks of relationships and the use of computational tools and interactive environments to help researchers ask new questions. Increasingly, digital libraries provide scholars with access to high-quality facsimiles of source documents in addition to (or instead of) carefully transcribed and marked up digital editions or structured databases.

Despite the rapid adoption of technology and its continuing impact on the ways in which cultural heritage material is presented and accessed, many aspects of humanities work practices remain unchanged. Digital libraries, with varying degrees of success, have achieved the goal of supporting tasks Unsworth calls "scholarly primitives:" discovering, annotating, comparing, referring, sampling, illustrating, and representing [2000]. These are primitives in the sense that they provide the building blocks scholars use to analyze, reflect on, interpret, and understand the material representations of the communities and cultures they study.

This dissertation follows the style of ACM *Transactions on Information Systems*.

These higher-order activities have proven more challenging to support. In describing when scholars should not use the TEI, Lavagnino notes, "the appropriate scholarly tools for the early exploratory stages of a project may be pen and paper, or chalk and a large blackboard, or a word processor" [2006]. These low-tech tools provide their users with the ability to quickly express and revise ideas without making premature commitments to a formal data model. They support the ambiguity and ease of use that is required for the creative exploration of ideas that is at the core of scholarship.

A preponderance of work within the interdisciplinary digital humanities (DH) field has focused on the use of technology to publish the "completed" results of scholarly research. Publication, however, is just the tip of the iceberg; the final physical (or digital) manifestation of years upon years of study. Digital libraries hold the potential to move beyond merely disseminating resources toward creating environments that support the analysis required to understand them. This dissertation looks at the rest of the iceberg: how can we design digital libraries that support scholars throughout the research process? Beyond publishing richly encoded documents complete with thoroughly vetted metadata or finding documents published by others, how can digital libraries support the iterative, open-ended, and idiosyncratic processes that scholars use to develop their own interpretive voice within a particular domain?

These questions are too big and offer too many fruitful lines of inquire to be addressed in a single dissertation. Within this broader context, I focus on how scholars use source documents (for example, early printed books, handwritten letters, medieval manuscripts, or modern authorial manuscripts) in the context of their work. To achieve

this, I conducted an in depth study of the research practices of humanities scholars. Building on these findings, I have designed and implemented CritSpace. Modeled in part on prior spatial hypertext systems, CritSpace designed to support analysis and critical engagement with source material during open-ended research tasks. It is implemented as an extensible framework for building highly-interactive, web-based, visual interfaces for content held in digital collections. These interfaces can be customized to support domain specific workflows that reflect the scholarly practices within different disciplines. I evaluated the usefulness of CritSpace by developing a customized implementation to support the ongoing work of a small research center.

1.1 Supporting Humanities Research

Research within the interdisciplinary digital humanities (DH) field has been dominated by the humanities community; often by individual scholars seeking to develop resources that meet the needs of their own projects or to embody theoretical visions of what digital scholarship should be (e.g., [DeRose et al. 1990]). This work tends to focus on describing the objects of study from within the framework of a specific theory, rather than the more traditional human-centered systems approach of analyzing the goals of specific user communities and the tasks they use to achieve those goals. The resulting tools may do an excellent job of supporting the humanities scholars' needs for "thick description" but often result in work practices that are intimidating to many scholars (for example, the expectation that scholars will manually encode documents

using XML or that emphasize topics such as authorship attribution that are far from the mainstream of humanities research).

Following the research culture of the humanities, digital humanists have often worked alone, producing individually "packaged" digital scholarly editions (for example, the Rossetti Archive [McGann 1996]). Medium sized projects like the Women Writers Project [WWP 2011] have emerged to create something akin to a monograph series; a series of digital editions produced to similar standard. On occasion, humanists have teamed up with information and library scientists to produce cultural heritage digital libraries that collect a wide range of textual and structured data (for example, dictionaries and catalogues of names) into a unified system. AusLit is an example of one such project [AusLit 2011].

In all of these diverse forms, technology has been particularly effective at providing new ways to record and publish the results of scholarly research. Moreover, once the sources for humanities scholarship have been recorded in some formal, machine readable format, computers can serve as incredibly powerful tools to help scholars perform automated processing—enabling them to harness this speed and power to pose questions that otherwise would be impossible due to the effort required to address them manually. Recent work has begun to emphasize the use of data-mining and information visualization to help researchers discover relationships (see, for example, [Clement 2009; Drucker 2004]).

These approaches, however, rely on some formal representation of the "data" to be used and the questions to be asked. Computers work well with structured information,

but, as Shipman and Marshall have noted, people are often unwilling or unable to describe that structure [Shipman and Marshall 1999b]. This holds especially true during the research process where scholars are working to gain an initial understanding of their source material. During this time they will develop provisional hypotheses that inform the short-term path of their research and will in turn be confirmed, revised and discarded in light of new discoveries. The structure that emerges from this process is vital to communicating their findings. By definition, however, that emerging structure cannot serve as the rigid *a priori* frame that software systems are inclined to provide. As in many other areas, scholars need tools that do not require them to commit to formal representations before they are prepared to do so [Shipman and McCall 1999].

Consequently, computer interfaces typically lack the simple power and flexibility of the most basic set of tools in the humanist's toolkit: pen and paper. At the same time, pen and paper have their drawbacks. Performing statistical analyses is rather more difficult and copying and pasting into your Microsoft Word document is out of the question. As my research has indicated [Audenaert and Furuta 2010], scholars frequently prefer electronic note taking. Notes on paper are most often useful only to the author himself, and sometimes not even to him. Scholars want their notes to be malleable: to add new observations over time, to share them with others, and copy them into new work. Sometimes paper is necessary, but it is rarely optimal.

The cultural heritage digital libraries in common use today have focused on publishing digital editions or on wide dissemination of source materials with few tools to aid in their use. To support these use cases, the humanities community has developed

detailed guidelines for producing, encoding and evaluating digital editions as scholarly resources. The digital libraries community has established standards for encoding metadata to ensure that the necessary descriptive and contextual information is recorded. In the best instances, these digital libraries are respected sources for difficult to access documents and serious scholarly editing. As Bradley and Vetch have noted, however, "digital libraries, as they currently are delivered, do not intersect terribly meaningfully with the process of scholarly research" [2007].

Digital libraries hold the potential to move beyond merely disseminating resources toward creating environments that support the analysis required to understand them. To achieve this, we must first develop a better understanding of how humanities scholars (and others) use source documents in their research.

1.2 Visually Complex Documents

Researchers in many different fields rely on digital facsimiles or other digital reproductions for many of their most time-intensive research activities. This freedom allows them to save precious travel time (and money) since they need physical access to the original artifact only to perform final proofing and verification work. In select cases, digital objects may fully satisfy the researcher's needs, thereby eliminating the need to access fragile originals.

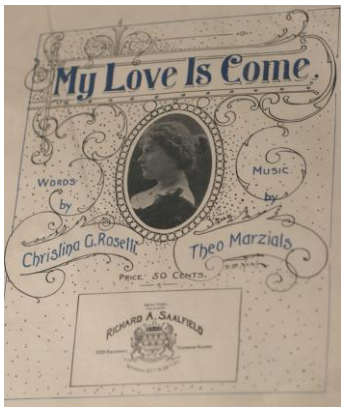

In this cover from the musical setting for a poem, the vignette of a woman suggests that the "My Love" referred to in the title should be understood as the woman in the vignette who is being addressed by the male speaker of the poem. Given the female author, this poem has traditionally been interpreted as having a female speaker.

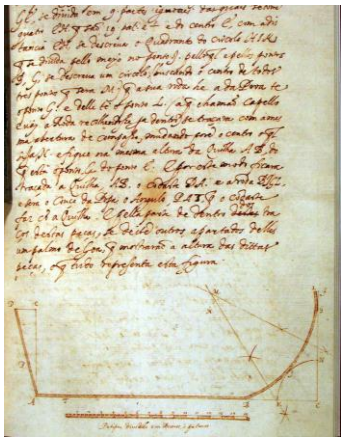

Ancient shipbuilding treatises contain both textual descriptions of the shipbuilding process along with drawings that demonstrate how shipbuilders performed their measurements. Nautical archaeologists look for information not only from the ink on the page but also by inspecting pin pricks that indicate where the authors placed the compass in order to make these drawings.

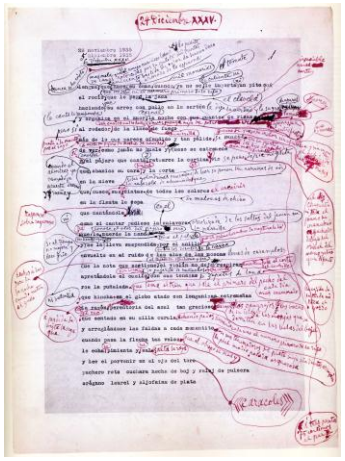

This authorial manuscript from Pablo Picasso (1881-1973) sheds light on his writing practice. Here he is working from a typed version of a previous (handwritten) draft of this work. A first set of annotations (additions to the text that will be added to the next typed draft) is written in black ink. The typed page has then been pasted onto a second sheet of paper and a second layer of annotations added in red.

Fig. 1: Examples of documents where the visual form of an object influences interpretation

Documents express information as a combination of written words, graphical elements, and the arrangement of these content objects in a particular media. The spatial arrangement and visual attributes (for example, color, font, size, and orientation) combine to form the visual grammar of a document or class of documents. This visual grammar provides components of the document's meaning in conjunction with the actual words or images of the document.

The relationship between the visual grammar of a document and its words and images varies as a continuum in (at least) two dimensions. On one axis, the visual grammar may be expressed more or less formally. For example, a journal article usually requires that documents follow specific layout guidelines, often specified using a set of styles in a Microsoft Word template. By contrast, the visual grammars employed in the documents in Figure 1 are unique to each particular document. Providing a complete description, let alone specification, would be difficult at best.

On the second axis, the visual grammar may be more or less integral to the meaning of the document. Changing the format of a journal article from that typically used by ACM conference proceedings to the format preferred by Springer's *Lecture Notes in Computer Science* is unlikely to cause us to think that these two documents are somehow fundamentally different. On the other hand, the visual grammar of an architectural design document is central to its meaning.

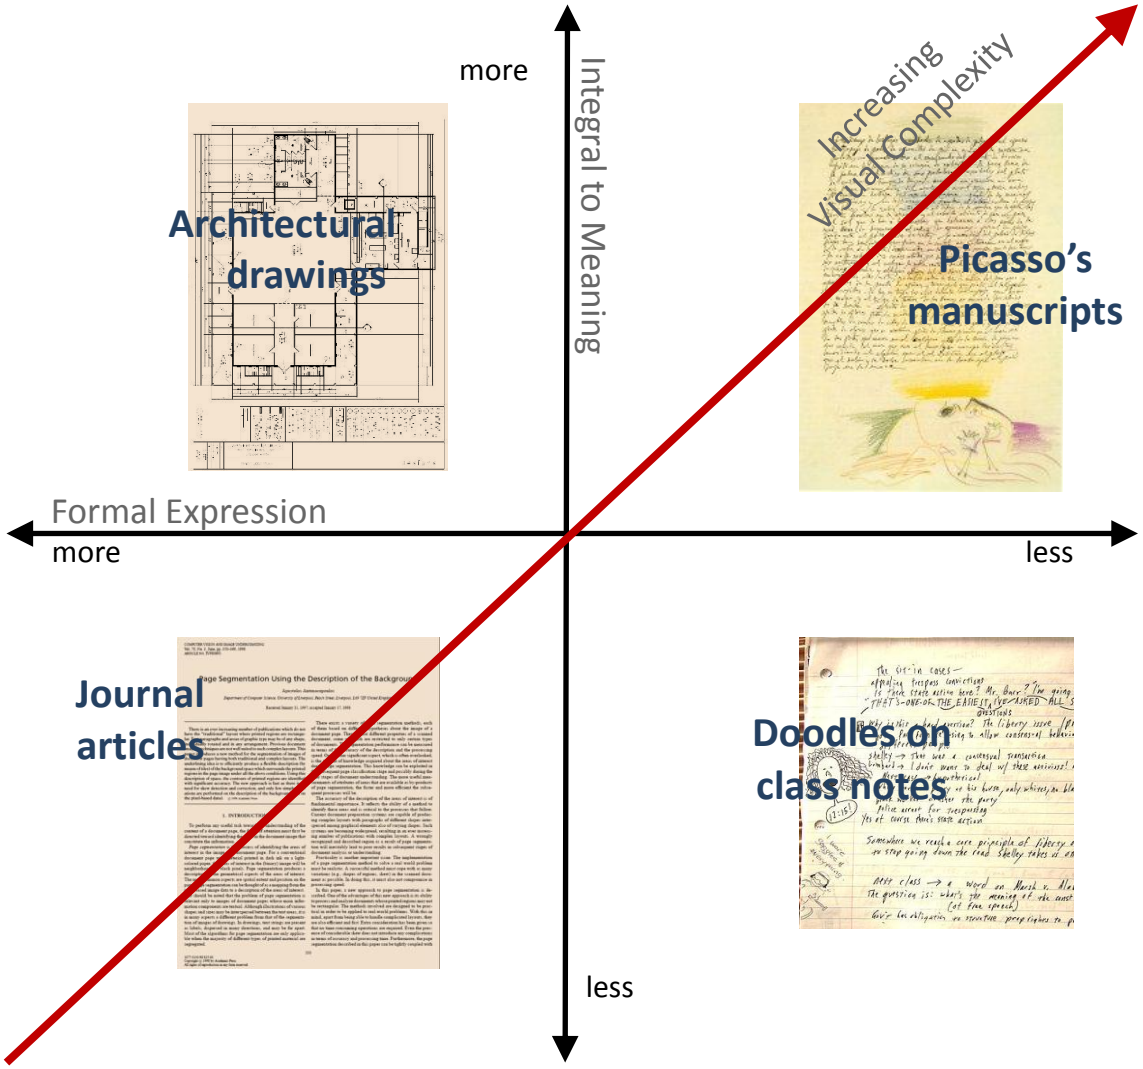

Fig. 2. Graph showing how the relationship between a document and its visual grammar varies along two dimensions

As illustrated in Figure 2, visual complexity increases as the visual grammar becomes less formal and more tightly integrated with the meaning of the text. It is important to note, though, that visual complexity does not depend only on the visual properties of the document. All documents convey information both in terms of the orthographic textual content as well as the way in which that textual content is arranged

on a page alongside (and structured by) other graphical features of the document. The importance of the actual words and images relative to the other visual content of a document depends on how that document is used. For example, the visual grammar of a printed book may not be particularly relevant when read as a medieval philosophical treatise. A scholar interested in the history of book making, however, may care more about the font size, illustrated letters, and layout of footnotes and scholia—the visual elements will take priority over the narrative content of the text. Consequently, the visual complexity of a document depends on a combination of the goals and objectives of the various people using the document and the properties of the document itself.

1.3 Approach

In this dissertation, I investigate the suitability of spatial hypertext as a basis for designing a creativity support environment (CSE) [Shneiderman 2007] to facilitate scholarly use of visually complex documents. Spatial hypertext systems provide an interactive visual interface in which users can gather and organize discrete information objects. These objects are represented as shapes in a two-and-a-half dimensional workspace and users can adjust visual properties such as the background color and border style. This simple, free-form approach to organizing information allows users to manipulate the layout and visual appearance of objects to informally express relationships and create complex organizations of an information space. Spatial hypertext has proven particularly useful for information triage (prioritizing an overabundance of data in time constrained conditions) [Marshall and Shipman 1997],

information seeking and structuring [Marshall et al. 1994][Buchanan et al. 2004], and incremental formalization of knowledge [Shipman et al. 1995].

Spatial hypertext systems enable lightweight interactions in the context of ambiguous or partially formed knowledge about a particular domain; knowledge that grows and transforms over the course of a user's exploration of the domain. This is precisely the setting of scholarly research. Moreover, the primacy of visual interactions with the visual source data of cultural heritage scholarship provides a promising approach to supporting interactive use of these materials.

On the other hand, spatial hypertext is heavily geared toward the organization of information, not using it. Existing systems typically represent content as small, atomic nodes of information, much like post it notes. This content can, typically, be opened in an external editor, but this disrupts the continuity of interactions within the work environment. Scholars will need better support for using content directly within the CSE. To help discover exactly what support they will need, Chapter II reports on a formative study in which I interviewed eight scholars from different disciplines in order to characterize why scholars use visually complex source material, what sorts of tasks they use these materials to accomplish and the likelihood that they will ultimately be willing to use digital tools in the context of their research. Building on the lessons learned from this study, in Chapter III I turn to review of prior work in spatial hypertext systems. This review looks at both how these tools seem well suited to supporting scholarship as well as the areas in which scholarly practice needs support that goes beyond currently available systems.

Chapter IV describes CritSpace, a web-based framework for building applications to support scholarly use of digital libraries. This chapter focuses on three key contributions of CritSpace:

- separating support for visual manipulation of content from the domain specific concerns of displaying and interacting with information
- supporting behavioral links between different content elements
- modular reuse and extension of system components

These three features of CritSpace are grounded in a framework-based approach to building spatial hypertext interfaces. The CritSpace system implements the core functionality of a general purpose spatial hypertext system, while custom extensions provide the behavior specific to individual domains and datasets that is required for intensive use by scholars. Since CritSpace must be deployed in the context of a specific application domain, I briefly introduce some of the features we have developed in support of textual criticism. Chapter V builds on this with a detailed case study that discusses how I designed and implemented the CritSpaceTC system for use in conducting a summative evaluation of CritSpace. Chapter VI wraps up with a discussion of the conclusions that can be drawn from this work and directions for future work.

CHAPTER II

WHAT HUMANISTS WANT

This chapter presents the results of a formative study designed to shed light on how scholars use source materials—original physical artifacts or their digital surrogates rather than transcriptions, scholarly editions or other derived forms of those objects—during their research. Before diving into the task of designing tools to support scholarly use of such material, it is necessary to first understand why scholars need access to visually complex documents in the first place and how they use them.

2.1 Prior Work

While it is commonly observed that there are few formal studies and little is known about how humanities scholars use digital sources of information (for example [Borgman 2009] and [Buchanan et al. 2005]), the library science community has produced a fair amount of helpful research aimed at developing services to meet the information seeking needs of humanities scholars. This work spans decades including the seminal work of Stone [1982] and Watson-Boone [1994]. Information seeking in this context covers a broad range of practices including but by no means limited to the use of electronic resources. Bates focused on scholars' use of a pre-web electronic system, Dialog [1996]. While this and other early work helps to characterize the information needs and behaviors of humanities scholars, the pace of changes in technology and

scholars' rapid adoption of these advances requires information professionals and system designers to continually re-assess current practice [Buchanan et al. 2005].

More recent work includes several studies on the strategies employed by historians [Dalton and Charnigo 2004][Tibbo 2002]. [Rimmer et al. 2008] discusses the implications of the physical nature of libraries on supporting scholarly work practices in digital ones. The LAIRAH project used deep-log analysis supplemented with questionnaires to investigate actual use of several UK-based resources [Warwick et al. 2008]. Whereas the majority of this research has relied quantitative measures (chiefly surveys) with large numbers of participants, Buchanan et al. adopt a qualitative approach toward understanding the information seeking techniques that scholars use [2005].

Ellis [1989], and later Chu [1999], develop models of the lifecycle of scholars' research process and argue that information seeking needs must be understood as evolving over the cycle of different research phases.

While much research remains to be done in this area, a useful portrait has begun to emerge. For a more complete survey of this prior work, interested readers are referred to [Toms and O'Brien 2008] and [Warwick et al. 2008] for useful, recent reviews.

Instead of examining information seeking behaviors, our work focuses on characterizing how scholars use the materials they find; how access to source material supports their research objectives, what types of insight this material provides and how creativity support environments might assist scholars' work practices beyond finding and retrieving information.

2.2 Methodology

Following a pilot study, I conducted a series of semi-structured interviews with eight participants (referenced as P1 through P8 in this description). With one exception, each interview took place in the participants' office or place of work and lasted approximately one hour. In the case of P8, the interview was conducted at a nearby Starbucks at the participant's request. Interviews were recorded, transcribed and analyzed. When relevant, the interviewer requested to see and photograph examples of source documents the participant was currently using.

The interviews consisted of open-ended questions that pursued a wide range of topics in order to encourage the participant to describe and give examples of his or her research objectives and the sources used in that research. Each interview began by asking the participant to describe their research area and to discuss what source documents he or she used. As the conversation progressed, the interviewer used a list of question topics to help direct the interview and ensure that the major points of interest were addressed. Answers to these questions were typically provided over the course of the interview without requiring the interviewer to ask them directly.

To select participants, I identified scholars (all faculty level researchers at academic institutions in the United States) who routinely use source documents in the course of their research. My primary goal was to obtain perspectives on how these materials were used across disciplines.

Accordingly, I selected participants from a broad cross-section of humanities research topics. Two participants (P4 and P7) represented the relatively common task of

scholarly editing. The critical bibliographer (P6) studied book production, systematically examining what the physical manifestations of an author's work may reveal about both the work itself and society's reception of it. The textual critic's (P8) work with New Testament manuscripts gave him a perspective on studies ranging from papyrus to medieval texts. The linguist's (P5) research centered on notes and official reports kept by colonial government and religious authorities.

P1 and P2 bring a perspective outside the mainline of humanities research. With P1, a laboratory scientist, I focused on the creation and use of laboratory notebooks. The research interests of P2, a nautical archaeologist, reflect the nature of his discipline at a crossroad between humanities research (examining old documents, tracing language change over time) and detailed observations characteristic of scientific field work. My interview with P2 focused on the creation and use of field notes and his use of historical documents as sources to inform his excavation and analysis of shipwrecks. P1 and P2 were included, in part, to help assess the degree to which the work practices we observe within the humanities can be generalized to research involving paper-based documents more broadly.

This diversity of participants provides a high-level perspective across a wide spectrum of humanities disciplines. The researchers I talked with had thought carefully (usually critically) about how others use source documents. Thus, while these conversations represent the objectives and peculiarities of individual scholars, their perspective on what is important and relevant to their scholarly community is indicative

of the broader needs of researchers than might otherwise be expected from so small a sample.

The reader should keep in mind, however, that our study objective—researchers working with source documents—represents a comparatively focused perspective on humanities scholarship. When asked about other scholars who might be interested in participating in the study, nearly every participant named one or two people. They were also quick to add that their research was not the mainstream, and that there were not many other people conducting similar research. The use of source materials remains an important but specialized task within the broader spectrum of humanities research. The results of our study, therefore, should not be interpreted as being indicative of all humanities research.

2.3 Findings

This section presents the findings from the study as they pertain to three key research questions: why do scholars bother with the time and effort to use source documents, what information do they hope to find and are they likely to use computers to support their research at all?

2.3.1 Why Bother?

Gaining access to source materials can be difficult and expensive:

> It's a kinda an ordeal to go to England, some of these libraries it's a pain in the neck to get into … manuscripts owned by private individuals, you know its like negotiating a foreign treaty to get access to see what they've got. (P4)

Once a scholar has gained access to the material, studying and analyzing the resources is an incredibly labor intensive process.

So why go to the trouble? Especially given the emphasis in digital humanities to prepare high-quality, richly encoded transcriptions, I wanted to know what advantage scholars found in having access to the originals (or digital proxies) of the source documents they used in their research.

2.3.1.1 Availability

The most obvious reason to prefer access to an original artifact to a transcription or other edited form is that, quite frequently, there is no suitable edited form. In order to create a critical edition a scholar must examine all of the source documents. A closely related task, creating a critical bibliography, requires the researcher to identify exemplars of all significant manifestations either of an individual work or, more commonly, of all works by a particular author.

This has significant implications for designing software systems to facilitate this research. While the production of scholarly editions comprises only a fraction of the total scholarly effort (one naturally expects more people to be using these editions than producing them), it is common enough that we had no trouble finding several scholars involved in this line of research.

These tasks are immensely time consuming:

> This project started in the mid-80s and we didn't publish our first volume until 1995. It took us ten years to get all of our stuff together and figure out what we were doing and so forth. And that year we published two volumes. And then it was five years before we published our third. It was another five

> years before we published our fourth. And this one will be close to another five before we publish our fifth. (P4)

P8, the textual critic, works with the more than 5,000 known copies of New Testament manuscripts:

> And so, our goal is to decipher the texts and transcribe them and if we had all the manuscripts digitally imaged right now it would take 400 man years to get the job accomplished.

He continues to observe that technology to speed this process would make "a remarkable difference." The long lives of these projects mean that technological innovation, as expensive as it may be relative to typical funding concerns in the humanities, may have a large impact in overall project costs. Measuring the opportunity costs related to funding technology investments, however, is usually beyond the expertise of project stakeholders.

While source materials are often used as the basis for creating scholarly editions, at times the contribution that a professional editor brings to an edition may not justify the required effort. Often, transcriptions of rare materials are designed to make those materials accessible to a broader range of scholars. In these instances, digital libraries that focus on delivering facsimile reproductions may be adequate:

> The other point is accessibility, and this is where historians win the argument, actually. They say the same question, why spend two years producing one monograph when we've looked at every single copy, right, when we can get grant money to digitize these archives and make them accessible. That's a fair point, it's a fair point. (P3)

2.3.1.2 Holistic Impressions

Contrary to common wisdom, scholars do judge a book by its appearance. Access to the original source often provides a high-level perspective that helps shape how a scholar understands a particular document's importance. P7, for example, described an alchemical notebook as follows:

> And, so what was striking about this was, you know, this is a, little quarto notebook that he used as like a laboratory journal to record experiments and things, and he did it with his wife, and it was somewhat random and haphazard in the first part of it and then she died. So he turns the book upside down and starts writing in it, … and the first four or five leaves are written in English, the rest is in Latin…

In publishing an edition of this work, the source material is not directly helpful to the intended audience, but the inclusion of a few images could help "to give someone a flavor, for what it was like" (P7).

Sometimes this holistic impression provides clues to a work's meaning that might be confirmed through subsequent investigation. These clues may open up interpretive possibilities that have been neglected by scholars. For instance, P6 describes the process of realizing that prevailing interpretations of a poem needed to be re-evaluated by examining a musical score:

> I know it doesn't look all that exciting, but, this is a poem by [a female poet] and although…you don't necessarily assume the speaker of the poem is the author, but [when] people read this romantic poem … they assume that the speaker is a woman. Well, when you put this picture of this cameo there, you, it kinda maybe shifts it. Maybe this is the beloved and not the speaker… and we played the music and it was like instantly, this was not written for a woman to sing.

While the holistic impression of a work may at times suggest an interpretation, at other times it commands an immediate reaction as in this excerpt later in our interview with P6:

> You can see what this morphs into. Stuff like this, see how much bigger this is than that. Um, [opening the book] oh, this one's crappy. There's not a place for writing in this.

While P6 and others commented explicitly about the importance of the visual form of a work, others tended to downplay its significance, treating the visual form as simply one possibly embodiment of the text. Despite their professed ambivalence toward the form of their sources, however, they frequently demonstrated a deep knowledge and attention to it that belied its role in their work. On describing why OCR wasn't possible for a set of town reports she had used years earlier, P5 observed:

> These are typed reports but … the typeset is bad, its inconsistent, it was real low end. It was a local press and the guy was just grabbing weird type set. Because like the 'e' would be here and then the next 'e' would be completely different typeset, I was like, he must have bought it at the used car lot or something…

Participants repeatedly provided colorful descriptions of their sources, even for documents they thought no one would be interested in seeing the original. The frequency and vividness of these descriptions suggests that an implicit knowledge of these texts as physical objects forms an integral part of their work regardless of whether that knowledge is ever described explicitly.

2.3.1.3 Nuanced Detail

Scholarship demands close attention to detail. Like crime scene investigators on a TV show, scholars scour their sources for clues that help them understand the text they are studying or that shed light on the social, historical, economic, political and linguistic context of its production.

P2 examines details in order to reverse engineer the geometry and other technology used to design ships, as evidenced in 16th century ship-building treatises:

> I mean its much better to look at an original, because you see the pressures, you see all the marks that will not come out in a digital reproduction. … you can see in the original, where did they put the compass. And the little holes are there… or the pencils that have been erased with bread.

P3 is trying to understand the geographic and political forces that may have influenced language change:

> So we take the time to do the paleography in the hopes of finding what we call "linguistic delicacies." That's what we get excited about. When did the change from 'f' to aspirated 'h' to silent 'h' complete itself?

Scholars sometimes find themselves fascinated by details in an effort to be systematic and thorough, without knowing in advance why a particular detail is relevant. P8 for example, was interested in new techniques to find out what kind of animal a parchment came from. When asked why this was relevant, he responded "um, I'm not sure exactly what the relevance is yet," and then speculated about possible implications:

> if for example, that this was, that they used antelope in a part of the world where antelope was not available, then that suggests either that the manuscript may have come from another place … So it opens up some possibilities of commerce and trade in the ancient world and things… And I think that can affect how rough that material is for writing on…might impact how the scribe writes.

Occasionally, the required level of detail or methodological protocol dictates that a scholar inspect the original object directly. Often, however, a high-quality digital facsimile is adequate:

> I'm pretty comfortable, when you take a really good scan like we've got of that St. Paul's [cathedral] manuscript, that I'm not going to see anything on the original that I can't see on that photograph, probably better. (P4)

While editors will try to identify and describe relevant details in their published editions, the level of detail required, the specificity required by different lines of research, and the need for visual inspection makes it impractical to describe all of this information in secondary sources. Consequently, many lines of inquiry require access to source material (either directly or through digital surrogates) even when high-quality editions are readily available.

2.3.1.4 Accuracy and Authenticity

Partially as a consequence of pursuing questions that hinge on nuanced detail, accuracy and authenticity are a chief concern among scholars.

While a scholarly edition may make a work more accessible, many scholars are reticent to trust the work of others if they have not been able to verify that information through personal inspection. When asked why scholars wanted to see the original rather than an edition, a pilot study participant replied succinctly, "because otherwise we would have to trust our colleagues."

P3 echoed this sentiment.

> The only reason I would look at a transcription that has already been done is to criticize it, to find fault in it. … No, I don't want to read someone else's

> transcription because I don't trust it, thank you. Particularly if it's a historian whose done it for his own purposes. Uh. Yeah, no, I'll make my own thank you.

These scholars are not just curmudgeonly, uncharitable sorts; they turn an intensely critical eye to their own work as well. For P2 the field notes are quickly transcribed and stop being used for day-to-day research needs "but you keep the originals because you make mistakes." For P7, ensuring the accuracy of his own work requires particularly careful attention when deciding how to transcribe punctuation. These smaller marks are more easily obscured by damage and may impact how a text is read:

> at the last minute that we're always going back and looking at these things and saying that really was, that's not what that says. And more importantly, or more typically, its, that really is a comma there, or a semi-colon, and you know it's the punctuation that's hard to see sometimes.

P3 explains the pursuit of seemingly minor detail in terms of methodological rigor:

> one might argue that the information that we provide is trivial, and sometimes it is: who cares if it's a cat-tail or una de gato is translated at this bush or that bush or whatever, you see. Well, so scientific inquiry dictates accuracy and confirmation on a fundamental level

2.3.1.5 Aesthetics

Less to the point of scholarly merit and rigor, the aesthetics of interacting with source documents was raised by nearly every participant.

> We're interested in the physical being of the artifact to the extent that we also show spines and covers, and end notes and end papers … you know I'm going to get a scratch and sniff screen one of these days to smell something musty and then it will really be like being in the library. (P4)

After talking at length about the need for transcriptions of New Testament manuscripts, the interviewer asked P8 about whether or not there was value in the original documents.

> There's always value in looking at the original manuscripts, otherwise, I wouldn't have a library like I've got, otherwise it would be all electronic books… But, the connection you make with a book when you read it manually...

He continued to supply a reason for that connection grounded the history of book technology. Contrary to the book being the quaint relic of yesteryear, "the fact that [a book] is a codex, while a computer is a scroll—scrolling is older technology."

2.3.2 What Is Their Quest?

The information we typically look to documents to provide—facts and figures, the literary content of a poem or personal account of a soldier—can easily be ascertained from a modern printed account of the work. What information, then, are scholars trying to discover by studying source documents.

While the specifics of what scholars look for depends on their particular research needs, we found four recurring themes in our interviews. Participants were interested in understanding the transmission of a text, surveying all known documents on a topic, identifying the agents who contributed to a text's creation and transmission and documenting the context (social, political, economic, intellectual, etc.) of its production. These four themes are by no means exhaustive but we found them to be particularly striking within the context of this study.

2.3.2.1 Textual Transmission

The most common goal in studying source documents that we encountered is the desire to understand how a text has changed over time. This is a special case of the theme discussed next, surveying all documents—specifically, surveying all documents that bear witness to an original witness.

In its most direct form, the objective is to identify the lost original or authorial form:

> we try to ascertain what, as closely as we can, which is an impossible goal, which is an impossible goal but we still strive for it, of what [the author] originally wrote. (P4)

Electronic variorum editions have been a common theme in digital humanities literature [Kochumman et al. 2002][Spencer et al. 2003]. In many cases, however, the model of document transmission as an original (lost) source that has been corrupted by errors introduced through repeated copying is too simplistic.

P3 describes the how records of early exploratory expeditions in the new world were created and transmitted.

> The chaps were required to keep detailed records and they did so in their own hand most of the time. And so we typically don't have those drafts… So they get back and they, through a scribe, re-write it. This already has major implications, right? … Same document copied five different times in five different places by five different people and with five different purposes. One's an informal summary, right, and omits things and adds things and condenses and combines and another one is a word by word and another one had three of four copies at hand and was intending to make the official… do you see?

In this case, recovering an accurate record of the original expedition may be an important objective, but not the only one. Understanding the transmission history helps

to illuminate what information was prioritized or reveal embarrassing facts that have been concealed. As a linguist, this participant was interested what these multiple copies might reveal about language use.

P6's interest in the transmission history has even less to do with discovering the original text. In the case of her author, the majority of her works were published within the author's lifetime. The authorial texts, therefore, are reasonably well-known. Her primary interest, instead, is on understanding the documents and how they shed clues on the author and her reception:

> to me its really interesting to understand a writer, not just, you know the way we look at her as English scholars is what shows up in the anthology, um, what poems did other writers get interested in, um, what, which of her books of poems seemed to sell the most, and, you know, important things...

In contrast to researchers who seek out the earliest, most reliable witnesses to an original text, this objective leads the researcher to

> trace the history of a particular poem as it emerged, say in a magazine, then in her collections then shows up in a hymnal and gets altered and then some…so you can really see how these things bounced around from genre to genre.

An authorial text is not the only form of content that is included in any given physical representation of that text. Tracing the supplementary material that is incorporated in the production process is another goal for some scholars:

> … when you look at the illustrations, you can pretty much tell that they were 19th century popular, some of it was popular religious art that I could actually trace… but, you look at these things, you know darn well that these were not specially commissioned for the book as this other person tried to say, its like, no, they went in the back, and they found what they had. (P6)

2.3.2.2 Survey of Evidence

A closely related goal is to survey all documentary evidence related to a particular research topic. P3, for example, described his research goals in the following way:

> the task is to look systematically at every primary document related to XYZ figure, XYZ topic, XYZ event. and pull out what's most interesting in terms of language, alright, and in terms of history, any number of perspectives are possible.

P2 expressed a similar need, saying that he wanted "to have a way for us to be able to retrieve, to say, give me everything on 17th century keels."

In many cases, modern editions may make this work more tractable, especially in the case of searchable digital texts. At other times, it is specifically the visual features of a text that are of interest. P8 compares how numbers were written in accounting documents as opposed to literary documents:

> the third feature that we have of the New Testament documents, they used the abbreviations for numbers …they typically would use the letters and they would put the bar over it… Well, uh, we don't see literary documents doing that at the time, we see accounting documents doing that.

While it is clearly feasible to record this information ahead of time if the encoder expects that it will be important, it is rare that the diversity of research needs that may benefit from particular set of documents can be anticipated fully. Thus, the systematic survey of a collection of source documents deserves careful consideration when designing cultural heritage digital libraries.

2.3.2.3 Agents

A fundamental task in understanding documents is understanding the people and institutions who produced them. An author's influence on a work is obvious, but it is by no means the only influence. Scribes, publishers, editors, illustrators, indeed even the audience, play important roles in understanding how and why a work came to take a particular form.

P8 describes how understanding a scribe's predilections can help determine the original wording of a manuscript:

> So for example if I have a scribe who is eliminating things that speak of Jesus's emotion…it probably would be a Gnostic scribe…but in that same manuscript I've got it saying …"Because He was angry, then He healed the leper", there's only one Greek manuscript that has that…I'm convinced now that the original text says that He was angry and that the others eliminated that because they didn't want to see Jesus as hotheaded

While P8, like most scholars working with medieval and earlier texts, infers a portrait of scribes primarily from internal evidence within a manuscript, P3 uses historical information and context to identify the individuals responsible for a particular copy:

> They'd be required to do a certified copy by the public scribe. We know these names for a lot of these time periods and regions.

Assessing the agents involved in a particular work is not restricted to looking at individual actors. In looking at how the Bible came to take the form it did among early Christians, P8 turns to a broad assessment of the scribal practices:

> I think we're dealing with some people who are not trained as scribes, as professional scribes, per se, who were the first scribes of the New Testament… but they were trained as accountants.

Similarly, P3 considers the religious affiliation and motivations of the scribes who produced the copies he studies:

> A couple Franciscans were very, very active in the 18th century, just going into the archives and rewriting sometimes summarizing, sometimes verbatim and this is a century after the events because they're religious folks required them or their own volition wanted to do this.

In looking at early printed materials, agency begins to take on a different importance. It becomes easier, for example, to examine the relationship between an author, people acting on her behalf and the publisher.

> …her literary executor, um, he, issued a couple of collections of her poems after she died, really important collections … and if he realized that he had messed something up, then he asked the publisher to correct it in the next edition or the next impression. I think I'm the only person on earth who knows this, because I've had access to [his] diaries and I can see, and because I've compared the books. (P6)

Scholars also pay careful attention to the material sources used in the production of a work for the insight that it may offer in understanding people and their motivations. For example, the bibliographer made an effort to identify the specific plates used to print different editions of the author she was studying:

> you can tell, you know, when the plates start getting banged up and, I mean, they start looking really cruddy. You know that they're just slinging the same things on there over and over and over again… you know, she was important but not that important (P6).

These material components, while not agents in their own right, embody the agency of publishers and offer valuable insight on the production and reception of the physical objects.

A frequently overlooked agent involved in the production of a work is the audience. The creators of a work are well aware of the tastes and preferences of their

audience. These preferences play a critical role in how authors, editors and publishers shape and market any given work.

> …it comes through on interlibrary loan and you know, the whole thing is printed in blue ink… Well, that would be a clue… or it [says] this appeals to the people who were buying this kind of book, which means her audience is skewed one way rather than another (P6).

An audience may also have a more direct impact. P6 relates how protestant editors repurposed a poem in ways that conflict with the original author's views and purposes:

> One of her hymns had two stanzas added to it, you know, and the point of those two stanzas seem to be to bring what she's saying in line with, um, like protestant ideas about sanctification at the time, which she as an Anglican would probably have some difficulties with, so, you know I see her work entering into all these other discourses and being circulated in these different ways.

2.3.2.4 Context

Texts and the physical documents they are written on exist within an embedding economic, social, political, and intellectual context. As hinted at in previous excerpts, understanding a document and its context are interrelated pursuits. P6, for example, discussed changes introduced by protestants to bring an Anglican's poetry into line with their notions of sanctification. P3 looked "systematically at every primary document related to XYZ figure, XYZ topic, XYZ event."

When asked about a particular document, participants would frequently start with a broad survey of the relevant context. P7, for instance, on introducing an electronic

archive he used, felt it important to provide a basic biographical sketch of the individual whose letters formed the basis of the archive:

> Hartlib was someone who served as a kind of, they described him as something like the Internet of the 17th century, he was of German/English parentage, and when the 30 year's war began, his mother was the daughter of an English merchant who was living in a Hunsiotic city in Northern Europe in Germany, actually today's Poland. And she married a German, sent her son to Emmanuel College, Cambridge, and then when the 30 years war broke out, he left Germany and went back to England and kinda lived there.

This sort of 'digression' was commonplace. In introducing any person or work, most participants felt compelled to introduce the interviewer to the appropriate context. This context, as above, involved people, dates, places, ancestry, education, political setting and more.

Not only is this information needed to properly understand the documents that participants were studying, but those documents also contain clues that inform unanswered questions about the context. A comment by P8 demonstrates this interchange of ideas:

> Some have suggested that the reason the Codex was used was because of the poverty of Christians, because so many of them were slaves, but even our early Codices on papyrus, seem to have generous marginal room on top and bottom and the sides so if its due to poverty, why did they leave so much room that way?

2.3.3 Will They Come?

Another issue that needed to be addressed before embarking on the labor intensive task of software development is whether scholars would be inclined to use digital study tools at all and, if so, to get a sense for the level of effort they might be

willing to invest. The contemplated creativity support environment would support information seeking and externalizing formative ideas.

While information seeking has received considerable attention, how to best support scholars' needs for externalizing knowledge is a question that has not been well studied. Much of the information that goes into the research process comes from a personal experience with a wide range of secondary literature and a deep, internalized understanding of the issues within a particular field. P7 hinted at this in describing the advice he would give to a new graduate student getting started on a particular topic:

> there is a vast body of secondary literature out there that you need to have some passing familiarity with…so if you're interested in [a poet] and the issue of women then there are certain things that you need to read just to, to get any sort of voice yourself, or perspective yourself, you need to understand the issue and its complexities,

This type of knowledge tends to be implicit and voluminous. Prior work has repeatedly demonstrated people are often unwilling or unable to express such knowledge formally [Shipman and Marshall 1999]. Any tools developed to support scholarly research must take this into account. They need to focus on helping the researcher to externalize knowledge as needed while recognizing that most of this information will never be formally expressed.

Given this prior work, I wanted to see if we could get a better feel for when scholars were inclined to record their thoughts, how formally or structured those records might be, what they chose to record, and why they recorded it. The participants in our study kept detailed, systematic notes (indicating a willingness to externalize knowledge) and predominately preferred to keep these notes electronically.

P7 described the importance of writing everything down:

> so, I learned a horrible lesson early on…I had a kinda breakthrough at one point, during my dissertation and I realized how all this stuff was fitting in together so I wrote it down, and later I was, re-reading these note cards and I had that, and I came across this idea and I said, my god, this is brilliant, and I'd completely forgotten it. You know it's what made my dissertation work, and made my career in a sense, and I had forgotten it… write everything down.

And P3 reports the level of detail in his notetaking:

> And I make my notes, line by line, because we go, literally, word by word, when you're comparing documents and line by line and all that.

These results should not be interpreted as contradicting previous findings about user's hesitancy to express themselves formally. I expect that participants tended to under-estimate (or at least under-report) the implicit, contextual knowledge that is latent in their notes. Instead, our findings indicate that as scholars conduct their research they are engaged in an ongoing process of incremental formalism [Shipman and McCall 1999] focused on a specific end product: a journal article or book. Their notes represent intermediate stages in this process. For example, P5 observed:

> then it's really easy because if I take all my notes that way and I type in the passages and put my comments under it, that's the start of my manuscript.

P3 reflected on the transient nature of the old hand-written notes he kept stored in boxes:

> No those notes turn into the commentary that's cross-referenced you see, so procedurally, those notes serve no purpose, except to form part of the famous [P3] archive of, which is never going to happen you understand. So they will serve literally no purpose except to keep some people warm in the winter if they ever discover them.

Significantly, despite the image of scholars as book-loving, technology fearing individuals, most of the participants preferred to keep their notes electronically. Even when they found making notes on a computer more difficult than writing on paper, the advantages for searching and for copying and pasting and for migrating into different forms (e.g., publication drafts) outweighed the drawbacks.

> Everything is electronic for me now. There was a fellow who did work on the pric adulterai, the story of the woman caught in adultery … it was 1700 manuscripts, something like that, all of his notes were on legal pads about this big of a stack. And you know where they are now? Still on those legal pads. So that was great research, why didn't you have the foresight to do this on a computer to start with. It's going to die with him… (P8)

Or P7 on finding that his computer battery died while researching in Europe:

> So that last week, I took notes by hand, and it seemed like it took me two weeks to transcribe those notes when I got back and I look on that as absolutely wasted time. Um, and, you know I would never take notes by hand, so I've always done everything electronically. Cause then it makes it a lot easier to incorporate into the written text.

Most participants recorded these notes using Microsoft Word or a similar commercial product. One used Zotero. These findings suggest both that scholars are willing to externalize more information than we originally suspected and that they could benefit from note-taking and annotation support build into their electronic work environment (provided that they can still export their data to more familiar products like Word).

Notetaking, however, is not merely the semi-formal representation of facts. Instead, it is an integrated part of the iterative writing process. Scholars think of their notes specifically in terms of how that information will be represented in the published form of their work and organize them accordingly. To be successful, support for note-

taking should be designed with this in mind and should provide a clear path to transition from informal notes to a final publishable manuscript.

2.4 Discussion

Access to source material is important. Digital facsimiles, while not fully adequate for all research tasks, play a critical role in mediating this access. In many cases, the digital copy of a document may be more useful than the original. At the same time, scholars tend to focus on the editorial contribution that they bring to a digital project. Their interest in computational support, at least within this study group, is in tools to help them prepare and disseminate their contribution to the scholarly community. They are less critically oriented toward their own work practices.

This lack of awareness of possible avenues for supporting humanities research with source documents, however, should not be taken to imply that no support is possible or needed. To the contrary, I would suggest that the clear (and relatively easy to achieve) benefits of applying technology to support the dissemination of scholarship, coupled with the comfortable familiarity of existing disciplinary methods has led the digital humanities community to overlook opportunities to critically assess how new technology might be developed to support the formative stages of scholarship.

Unsurprisingly, but easily ignored by non-specialists and taken for granted by specialists, the use of source materials in scholarship is not a simple, straight-forward examination of an artifact in isolation. Quite to the contrary, for participants in this

study, the "primary" materials being considered seemed to occupy a comparatively small degree of attention—even in cases where they were the focus of the research.

Instead, source materials are studied in order to address broader questions: how was this author's work received, what are the economic structures that enabled this activity, how is language change being influenced by the social-cultural forces of a particular geographic region? Even for scholars working on traditional textual editing tasks, the goal of reconstructing the "original" version of an author's text transcends understanding a single document. An individual document in this setting may be prized, not because it sheds light directly on the content of the original, but because it informs the editor about the habits of the community responsible for transmitting the text.

Scholars study source materials as an integral part of a complex ecosystem of inquiry that seeks to understand both the text being studied and the context in which that text was created, transmitted and used. Both the content of the document and the embedding context are known in part and scholars seek to systematically examine the current body of knowledge in order to develop a more complete understanding of this entire ecosystem. This study proceeds in a dialectical fashion, with insights gleaned from one document informing the knowledge of the field which, in turn, invites further examination of the original source and related documents.

2.4.1 Supporting Use

The cultural heritage digital libraries are seen primarily as tools to enable the digital publication of humanities scholarship. These libraries provide wider access to

source material and to the editorial contributions that have been invested in their creation. They also provide scholars with tools to document difficult material that cannot be effectively (or affordably) disseminated in print form. Digital libraries of varying levels of sophistication have become standard tools of practice across the humanities disciplines for access source documents.

However, these resources are conceived of primarily as final repositories for scholarly contribution. We submit that there is a pressing need to rethink this conceptualization. Digital libraries hold the potential to serve as both valuable resources for finding and disseminating the results of scholarship as well as sites to support research in process. In particular, given the size and scale of many projects that focus on source material—projects that often measure their duration in decades rather than years—we need to consider the role of digital libraries as evolving resources designed to support the entire life cycle of a research project. This life cycle should include the initial digitization of material, the ongoing analysis of that material and the ultimate publication of long-lived scholarly resources.

We highlight five important implications from this vision of digital libraries designed for the humanities scholars:

First, the environment that supports scholarly work is as important as the mechanisms for information retrieval and the metadata. Designing, developing and maintaining these environments will require a degree of ongoing technical investment that is rarely found within the humanities community.

Second, these libraries will be focused. Instead of millions of documents, most humanities projects will have a few thousand items at most. Collections numbering in the hundreds will be the most common. These items, however, will have complex internal structure and will need to be supported with access to a wide variety of related contextual material and editorial notes.

Third, as a site for ongoing scholarship, these libraries become bootstrapping tools for their own construction. That is, a library may be developed initially, not as a resource for a general audience, but as a repository of digital facsimiles. These facsimiles will then be used internally to support the work of an individual editor or editorial team whose output will be integrated into the original library that will be of use to a broader audience. This bootstrapping process, not just the final content or form of the library, needs to be carefully considered in the design process.

Fourth, the long life-cycle of projects in the humanities—both the development phase and the subsequent use of the resource—means that these libraries will be developed as an ongoing process with changing audiences and needs. The architecture of these systems should support changing needs for both data representation and user interface components.

Fifth, designing and maintaining libraries with this degree of complexity requires a considerable investment. This investment, while not justified in all instances, has the potential to yield compelling returns when considered in the context of projects that frequently last more than a decade. Far from being uncommon, P2, P4, P6, P7, and P8 were all involved in work of this duration or longer.

CHAPTER III

ASSESSING SPATIAL HYPERTEXT

In designing support for scholarly practice in the humanities, I have focused on using a highly interactive, visual interface that is heavily motivated by prior work in spatial hypertext. This approach to interface design offers proven techniques for supporting emerging knowledge structures that are central to creative work such as scholarly research. At the same time, my research contributes to a deeper understanding of how to extend current models for spatial hypertext to help scholars more effectively use the objects they find in cultural heritage digital libraries and archives. In this chapter, I briefly introduce spatial hypertext[1] and then discuss its relevance for supporting creativity in general and humanities scholarship in particular.

Spatial hypertext emerged in the early 90's from prior experience with map-based hypertext system such as Trigg's tabletops [1988] and Conklin's gIBIS argumentation support system [1988] (shown in Figure 3). In contrast to navigational hypertext, in which links between documents or information nodes are used primarily for navigating a network of information (like the Web today), map-based hypertext provides a top-down view of that network, allowing people to see and manipulate structured relationships between different content. Aquanet [Marshall et al. 1991] built on this work to allow users to define their own formal schemas to represent relationships

[1] For a more complete survey of this field, please see [Shipman and Marshall 1999a], [Shipman et al. 2001a], and [Bernstein 2011].

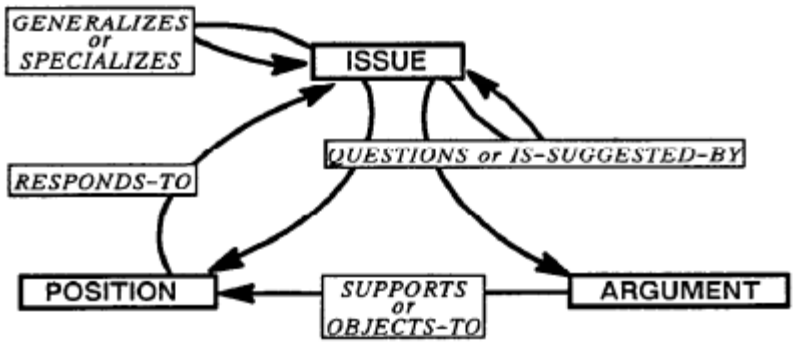

Fig. 3. The structure of an argument in gIBIS demonstrating the use of map-based hypertext to specify formal structures and relationship between information nodes [Conklin 1988].

between different content objects (typically user-authored notes or links to files). One of the core strengths of these systems is that they allowed users to work with information visually; enabling them to express concepts and formal relationships and to see those structures within large collections of "information nodes." Once expressed formally, these knowledge structures can be used to inform computational analysis.

In deploying and evaluating Aquanet, however, [Marshall and Rogers 1992] found that users were often unable or unwilling to express relationships between nodes formally. Early in an information analysis task, they did not feel like they had a strong enough grasp of the problem domain to commit to a formal representation of their knowledge. As they became more confident in their understanding of a domain, a new problem arose: the task of explicitly expressing their knowledge interrupts their analysis work and confers limited advantages. The benefits of representing relationships between information explicitly did not justify the required effort.

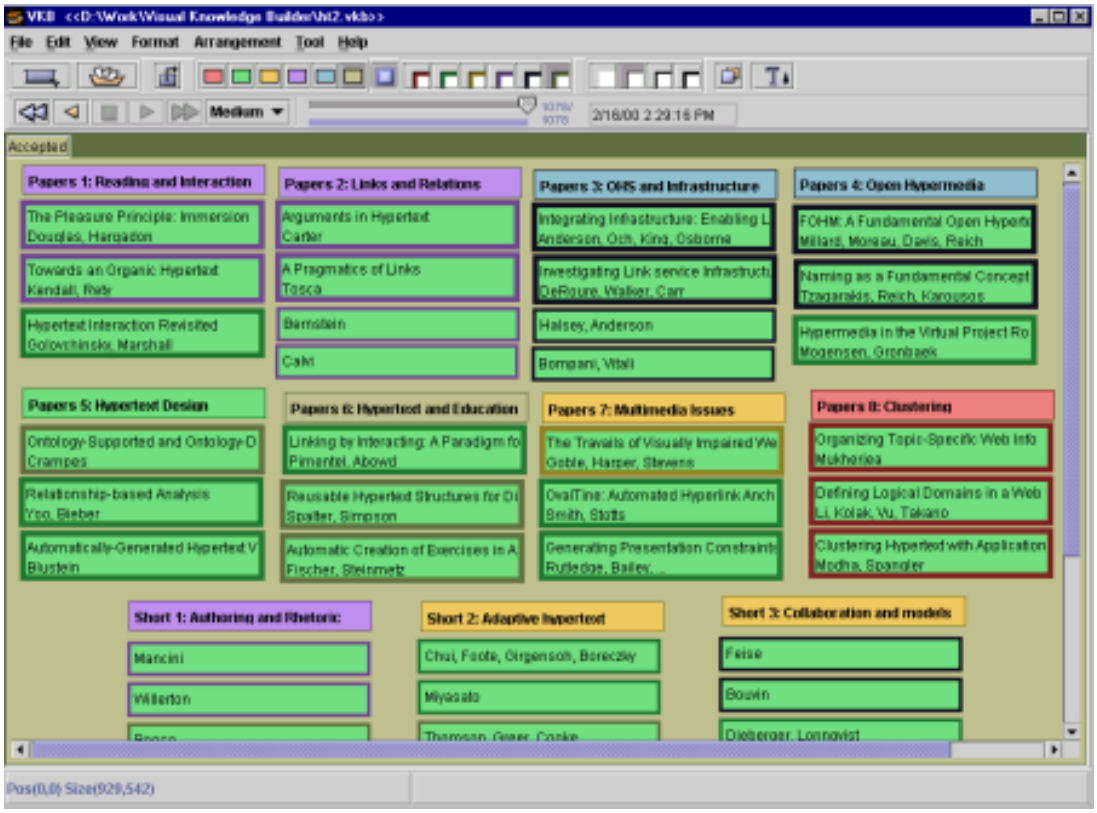

Fig. 4. A spatial hypertext used to organize talks at a conference. From [Shipman et al. 2001b]

While Aquanet's users did not want to be forced into premature commitments about structure, they did use the system to provide informal structures, by grouping related items into stacks and piles and by changing other visual properties. VIKI [Marshall et al. 1994] and later VKB [Shipman et al. 2001a] arose from these experiences in order to support interaction with these implicit structures. As shown in Figure 4, these systems allow users to represent relationships in "spatial hypertext" by manipulating the visual properties of content objects. Typical visual properties include spatial positioning in a two-and-a-half dimensional workspace, font style, border color and thickness, and background color. The addition of a spatial parsing [Marshall and

Shipman 1993; Shipman et al. 1995] enables the system to recognize the implicit structure represented by the hypertext and support user interactions with that structure.

Spatial hypertext principles have motivated tools designed to support a variety of different tasks including organizing personal collections [Shipman et al. 2004], digital library interface [Buchanan 2004], and authoring literary hypertexts [Bernstein 2002].

3.1 Supporting Scholarship

Spatial hypertext systems provide support for lightweight interactions that aid users in gathering and organizing data during a variety of information tasks. They are designed to minimize interference with the user's work-process, what Csikszentmihalyi calls the "Flow" of creative work [1990]. Shipman and Marshall highlight four key features that make a spatial hypertext system useful in comparison to a document-centric interface [1999]:

1. It takes advantage of people's considerable visual recognition and intelligence
2. It facilitates constructive ambiguity
3. It supports emerging problem solving strategies
4. It reduces overhead in communicating with others

Each of these benefits represents a key need of tools designed to support scholarly use of source material in cultural heritage research. The first benefit stands out in particular. Spatial hypertext's emphasis on visual representations seems particularly

well suited to helping scholars interact with visual representations of source documents and other image-based representations of cultural heritage objects.

Consequently, spatial hypertext provides a promising starting point for supporting scholarship. In addition to these general benefits, I want to highlight two systems that have approached spatial hypertext in ways that are particularly relevant.

3.1.1 Representational Talkback

[Nakakoji et al. 2000] calls attention to spatial hypertext as a paradigm for building interactive systems that support the early stages of design. They define design quite broadly, emphasizing writing as a design process, but including all activities that are characterized by the interdependence of problem analysis and solution synthesis. Following [Schoen 1983], they postulate that the design process involves a designer creating external representations of his analysis and/or solution. This initiates a "reflective conversation" with these externalizations in which the material representations talk back to the designer and thereby inform subsequent stages of the process.

Thus, design proceeds as "the designer acts and reflects almost simultaneously; acting, interpreting and reacting to the evolving design," a process Schoen calls "reflection-in-action." Reflection-*in*-action stands in contrast to reflection-*on*-action, in which the designer pauses in his action, interrupting a state of focused concentration on a task in order to reflect intentionally on the external representation of the problem. While both types of reflection are important in the design process, reflection-*on*-action tends to

be more important in later stages while reflection-*in*-action is the preferred mode early in the process.

Observing that traditional computer applications support workflows that lend themselves more to reflection-*on*-action, Nakakoji et al. propose two dimensional spatial positioning as a strategy for supporting early stage design work [2000]. Their ART (for Amplifying Representational Talkback) system allows writers to begin authoring chunks of text that are displayed in small rectangular blocks. These blocks can be freely arranged in a two dimensional workspace. The vertical position of these blocks is used to generate an ordinary linear view of the document in progress. [Yamamoto, et al. 2002] later extend this work to other authoring tasks including note-summarization, multimedia data analysis and movie editing.

Based on user-studies for these systems, Nakakoji et al. find that their approach does in fact support the early stages of authoring by allowing users to externalize thoughts without requiring commitment. They emphasize the creation of domain-specific solutions in order to ensure that the interactions during the design process minimize the cognitive effort required by the designer. In this work, spatial hypertext serves not as the primary content (that role belongs to the linear document being authored) but rather as a medium for interacting with material externalizations in a way that facilitates reflection in the midst of the design process.

3.1.2 Notetaking and Annotation

Lamenting that "digital libraries, as they currently are delivered, do not intersect terribly meaningfully with the process of scholarly research," [Bradley and Vetch 2007] describes Online Chopin Variorum Editor (OCVE). Following [Brockman et al. 2001], they identify the twin tasks of note taking and annotation[2] as being of central importance to the practice of scholarship in the humanities. The OCVE prototype developed for this project was designed to help users understand the relationship between multiple witnesses to a single work by Chopin. While not self-consciously a spatial hypertext tool, OCVE provides an interactive visual interface that supports superimposition, juxtaposition and combination/interpolation of images of sheet music, in addition to tools for annotation.

Annotations were not displayed in the same visual workspace used to display the images of sheet music. Instead, they were attached to the images and available for viewing on demand in a dialog window. In practice, the authors found that this approach to annotation yielded mixed results. While OCVE was adequate for viewing the annotation, creating annotations was problematic. This, they found, was a "fairly profound problem" since, in stark contrast to making notes on paper, the interacting with the software was perceived by users as a hindrance to recording responses. They note that "the annotation process is highly sensitive to the practical burden it places on the mind of the annotator."

---

[2] Note taking involves jotting down thoughts, comments, observations and questions that may or may not be tied to a specific source material or that may refer (explicitly or implicitly) to multiple sources. Annotations, by contrast, are notes or other markings that explicitly reference an object of study.

In separate work, [Bradley 2008] introduced Pliny as a thought-piece designed to focus specifically on the task of supporting scholarly annotation. In this tool, he explicitly invokes spatial hypertext as a design paradigm. A primary goal of Pliny is to support the ordinary work practice of scholars, introducing 3x5-card-like note taking behavior in a medium that is more pliable than paper. Pliny allows users to attach note-card-like annotations to images, web pages and other objects. It also supports a hierarchical model of note-taking, in which a note is paired with a 2D display area that can contain multiple other notes in a VIKI-like display. Containment structures enable the user to group related notes. The display objects in Pliny are references to the source note, thereby allowing a single note to be added to different display screens.

Unfortunately, Pliny is a desktop application based on the Eclipse software development platform. The complexity of this tool combined with its status as a thought-piece has prevented its widespread adoption by the user community it purports to support. Additionally, Bradley has not published information on user evaluations of Pliny. While Pliny appears to use spatial hypertext to support humanities scholarship in innovated ways, it remains difficult to assess the ultimate usefulness of this system. Hence, Pliny remains a promising but untested model for providing scholarly support. Commercial systems such as One Note [2011] and MindManager [2011] that adopt many of the features of spatial hypertext systems provide more polished user interfaces. Even so, judging from the limited scope of the formative study, these tools have not yet had a significant impact on the practice of scholarship. Further investigation is needed

to determine the level of awareness of these tools and to ascertain what, if anything, prevents their widespread adoption within the scholarly community.

Despite these experiences, there is a pressing need to support note-taking and annotation and spatial hypertext offers a promising model. When asked about their use of computers, the participants in the formative study expressed a strong preference for taking research notes on computers. This wasn't always convenient (for example, when the nautical archaeologist needed to make sketches of physical objects). Much like OCVE, when pressed, participants reported that electronic note taking required intentional actions that interrupted their work process. Despite these interruptions, however, many found the difficulties worthwhile because electronic note taking conferred two key advantages: participants could find their notes more easily, and they could copy and paste their notes into the draft manuscript they were working on. Supporting these key requirements, while simultaneously minimizing interruptions to the scholar's workflow, is one of the core competencies of spatial hypertext systems.

3.2 Does the Shoe Fit?

The motivation for this research is to identify ways to support creative research involving cultural heritage data in general and source documents in particular. Spatial hypertext systems offer a proven approach to supporting creativity in exploring complex, poorly understood information spaces. But do spatial hypertext systems really match the working needs of humanities scholars and others working with cultural heritage resources?

In the formative study, I identified two work practices that stand out as representative: gaining a high-level, visual overview of a physical object and detailed analysis of an object. Spatial hypertext seems particularly well-suited for the first of these tasks. It is straightforward to see how this general approach to gathering and organizing information can be used to allow readers to perform a quick visual inspection of an object. By retaining that object in the workspace and allowing the user to position it relative to other objects, the system leverages human's ability to quickly process visual and spatial information. Thus the object comes to serve as a bookmark that references certain ideas and situates those ideas in relationship to other bits of the intellectual puzzle the scholar is pursuing. This meshes well with the notion that spatial hypertext promotes creativity through reflection-in-action [Nakakoji et al. 2000].

The second of these work practices, detailed analysis, presents challenges. Spatial hypertext systems have been designed to support information retrieval and organization. While this involves analysis and reflection on an information space, that reflection is usually mental—focused on the ideas and concepts represented by the nodes displayed in the system. In the case of one major application area, information triage, this work is performed under stiff time pressure. The challenge is to quickly refine from a superfluity of data to prioritize the information that is relevant to a particular task. This represents a domain mismatch from the systematic, exhaustive study of a body of information and the close inspection of nuanced detail that is characteristic of humanities research. While time-pressure is inescapable in every discipline, humanities scholars are

more likely to measure their research projects in terms of decades or lifetimes rather than hours or days.

The content of traditional interactive visual interfaces typically consists of small, user authored text blocks, document proxies returned by a search engine (with a link to an external application to open and read that content), images, music files, or video along with supplementary metadata. Spatial hypertext systems tend to treat this content as atomic blocks of information and provide little if any support to manipulate internal structure. The information that humanities scholars use, by contrast, may have rich internal structure and require interaction with that structure. Digital facsimiles have pages and chapters, to pick a simple example. While users may be unwilling to express the formal relationships between objects, the contents of cultural heritage digital libraries contain rich, well-defined relationships and scholars think of this material in terms of these relationships. To support cultural heritage scholarship, researchers will need interfaces that allow for explicit representation of these disciplinary domain models.

It is important to note that this does not negate prior work about the need for ambiguous knowledge representation, avoiding premature formalization of knowledge or people's unwillingness to break out of the flow of their work to make the formal relationships between information nodes explicit. Quite to the contrary, scholars need to focus on the task at hand and avoid the distractions of working with meta-information about that task as much as anyone. Instead, support for humanities scholarships requires systems that enable built-in support for pre-existing forms of structured information and

that present this information in ways that reinforce the fluid manipulation and informal interactions characteristic of spatial hypertext environments.

In summary, spatial hypertext holds promise for supporting humanities scholarship, especially for work that requires access to image based information such as source documents. The ability to support open ended exploration of ideas and to represent relationships between visual materials visually fits well with the tasks of cultural heritage scholarship. But there is a domain mismatch between the information seeking and structuring tasks spatial hypertext systems have traditionally been designed to support and the in-depth analysis required by scholars. To support scholarly analysis, we need to introduce features into spatial hypertext that can take advantage of existing structure and support domain specific interaction with objects in the workspace. In the following chapter, I describe a new system designed to do just that.

CHAPTER IV

CRITSPACE

CritSpace is a web-based creativity support environment designed to provide a user interface to cultural heritage digital libraries that supports interactive use and critical analysis of digital content. A key design goal of CritSpace was to enable both high-level information gathering and structuring tasks and careful analysis of objects in the same visual workspace. It focuses on supporting the use of source material, that is, image-based representations of material objects, rather than secondary material such as scholarly analysis of those objects.

By supporting creativity, I mean supporting the open-ended exploration of a collection of material that is intended to result in a better understanding of that material or a new scholarly perspective. This includes a variety of work-process stages such as fact-gathering, where the scholar is encountering and analyzing objects; idea generation, where he begins to articulate a unique perspective or "scholarly voice" on that material; and a validation stage in which he tests emerging ideas against the material facts to confirm hypotheses and further inform the creative process. Far from being distinct, sequential phases, these processes are envisioned as analytical descriptions of work that is frequently conducted in a deeply inter-twined fashion, often with all steps happening simultaneously.

To achieve this support, CritSpace builds on the foundation laid by previous work in spatial hypertext, while introducing features that facilitate close, domain specific

interaction with the materials displayed in the workspace. As previously noted, in designing CritSpace, I have emphasized three major contributions:

- separating support for visual manipulation of content from the domain specific concerns of displaying and interacting with specific types of information
- supporting behavioral links between different content elements
- modular reuse and extension of system components

Instead of trying to create a general-purpose application to meet all scholarly needs, CritSpace is a framework for constructing web-based, visual interfaces for cultural heritage digital libraries. As such, it relegates the details of providing access to specific content and support for domain specific user interaction to implementations of the framework. CritSpace itself provides a core implementation of visual objects (called Panels) that allow users to manipulate content in a two and half-dimensional workspace. The core system also supports transparent persistence of the current state of the workspace. Figure 5 shows an initial prototype of CritSpace used in conjunction with the paintings and writings of Pablo Picasso.

The CritSpace framework is implemented (or deployed) in a particular context, by creating an initial web page that allows users to load existing workspaces or create new ones and then by defining a set of Panel extensions to be used in the workspace. Implementing a trivial instance of CritSpace, one that allows users to enter notes in boxes and organize that information in the workspace, is a straightforward process requiring no more than a few dozen lines of code. Such an implementation allows users

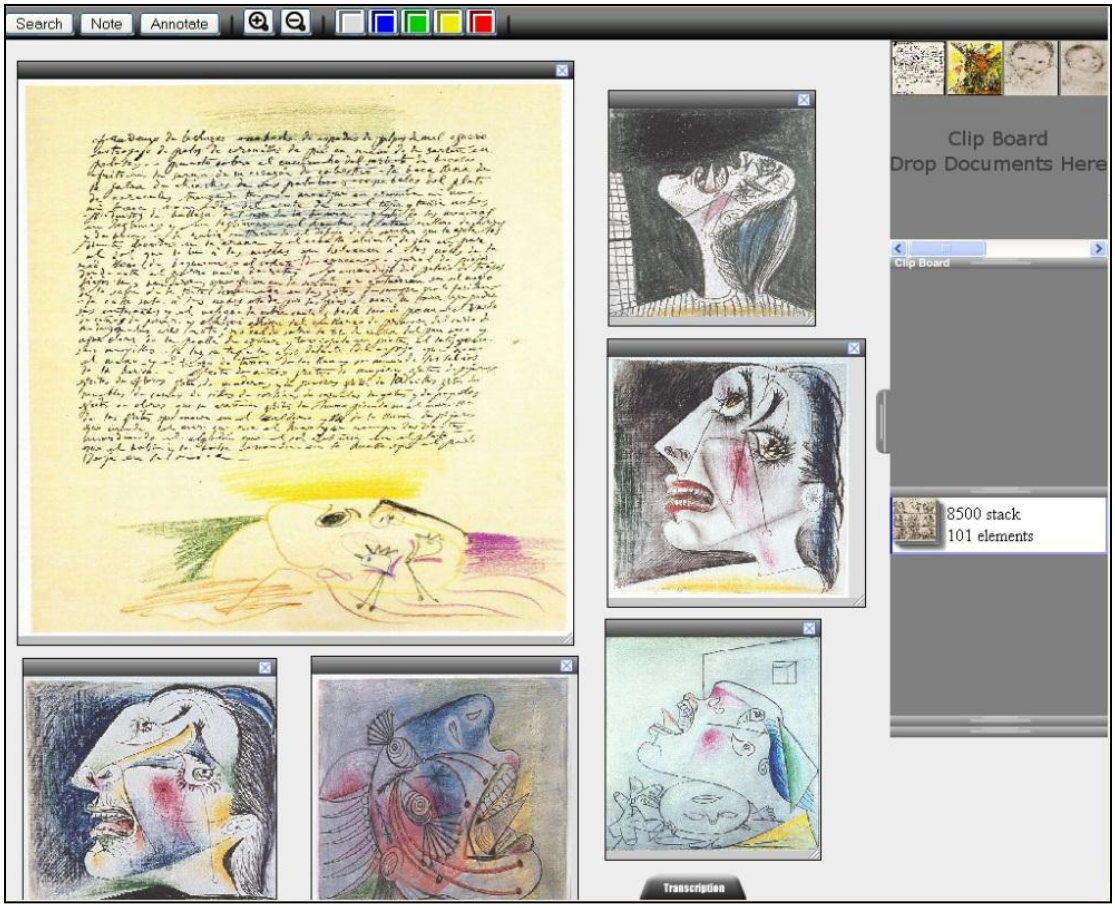

Fig. 5. An early version of CritSpace showing a poem by Pablo Picasso along with preparatory sketches for his artwork Guernica. This juxtaposition highlights the similarity between the colors and the figures that are present in both the written work and his sketches. From [Audenaert, et al 2008].

to create notes, arrange them in the workspace and adjust their visual properties such as background color and border style.

Naturally, most useful implementations will require more work. As a prototype system, I created a CritSpace deployment designed to support research in the field of textual criticism called CritSpaceTC. This system allows a scholar to compare a base-text with a high-resolution digital facsimile in order to identify all places where a

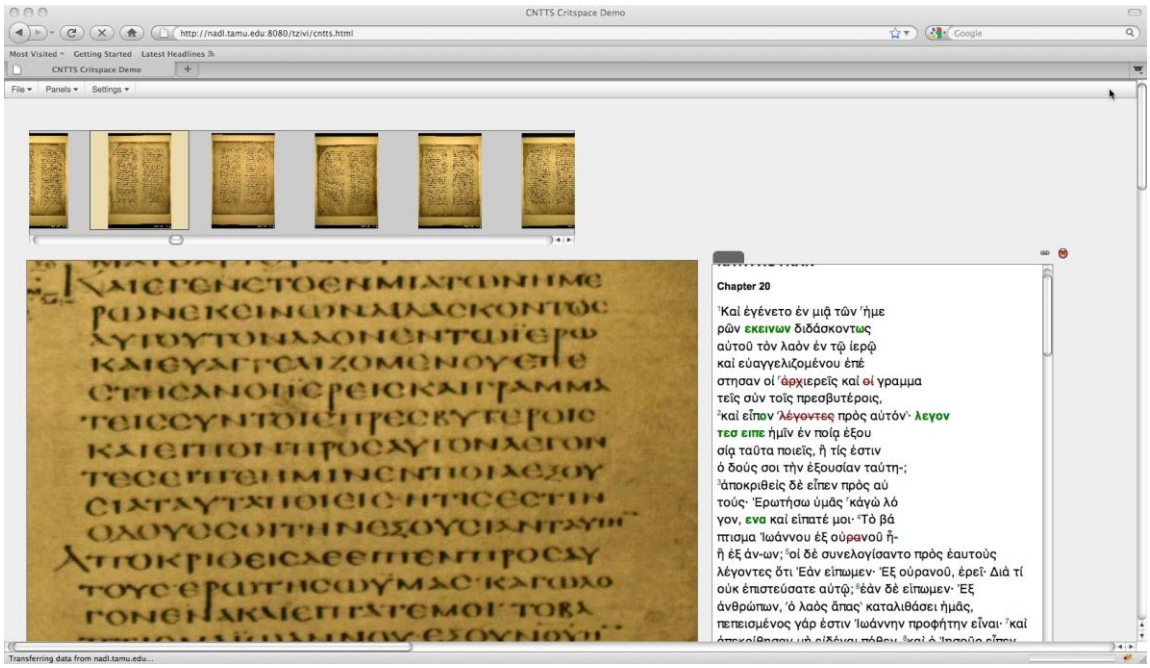

Fig. 6. Screen capture of a user interaction with CritSpaceTC

manuscript copy differs from the base text. While this particular application domain is not new, it does serve to demonstrate some of the key features of the CritSpace framework, including the use of complex user interactions within the content of a Panel, support for both system and user initiated links between panels, and integration of separately developed web services. Notably, the textual criticism domain is a far cry from the information seeking and information organization tasks that have motivated much of the initial research in spatial hypertext systems. In this instantiation of CritSpace, the focus is placed squarely on discipline-specific use of material retrieved from a digital library over gathering and organizing information.

Figure 6 shows a screen capture at the end of a typical session of a user study we conducted with the system. CritSpaceTC displays a digital facsimile in a filmstrip style

panel that allows users to quickly browse the entire document. Individual pages can be dragged from this main display into the workspace, creating a PageDisplayPanel that allows users to interact with high-resolution images using zoomable image viewer. This panel is synchronized with the main facsimile display so that if the user selects a new page in the main facsimile panel, the page display is updated simultaneously. A BaseText panel is used to display, in this case, the theoretical "original" text that will be compared with the manuscript to identify variant readings. It uses a heavily customized form of the Rich Text Editor widget from the YUI library that lets a textual critic to mark variants on the base text [YUI 2011]. This panel can be explicitly linked to the PageDisplayPanel. Once linked, the BaseText panel adds a new overlay layer into the image viewer that allows the user to place verse markers on the manuscript images (for Greek New Testament manuscripts). These markers are linked to the base text so that when a user clicks on a marker, the corresponding verse is highlighted and scrolled to the top of the panel's display area.

The remainder of this chapter describes the CritSpace system in more detail, the following chapter builds on this discussion to describe the system in a more human-centered context based on a summative user study.

4.1 The Framework Scaffolding

The bulk of CritSpace's power comes from the Panels and PanelLinks that people use to interact with digital objects and from the visual properties module that implements the spatial hypertext features for organizing that content. These components

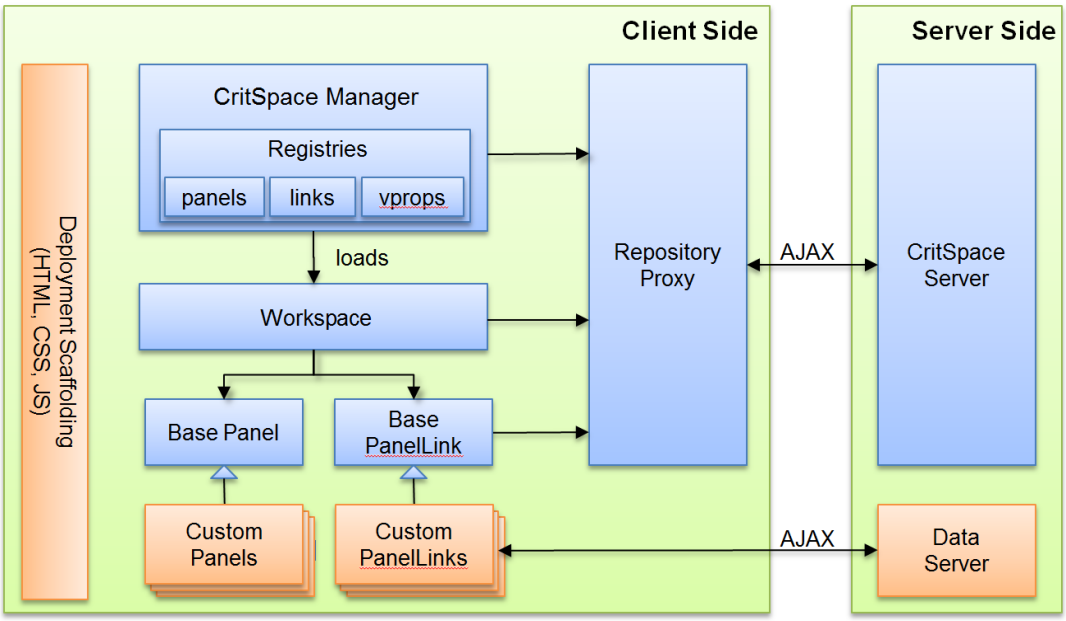

Fig. 7. CritSpace system architecture

are discussed in detail in later sections of this chapter. First, we take a bird's eye view of the scaffolding that ties these components into a flexible framework and makes it suitable for use in many different application domains.

4.1.1 System Architecture

Figure 7 shows a high-level view of the CritSpace architecture. Components that are part of the CritSpace framework itself are shown in blue, while deployment-specific customizations are shown in orange.

The main CritSpace server provides a web-based API to the data persistence layer. The majority of the system is implemented in JavaScript that runs in the user's web-browser. The client side components access this server through a pluggable repository object. By encapsulating the details of network communication, the bulk of the CritSpace framework can be implemented without worrying about the complexities of asynchronous communication with a remote server. This encapsulation also provides a layer of abstraction, allowing significant flexibility for later changes to the server implementation to have little or no impact on the main codebase of the client.

The starting point for the client-side application is the CritSpace Manager. This object holds the primary context for loading the repository proxy that manages the connection to the server as well as the registries for panel extensions, panel links and visual properties (details on these below). In addition to providing the context for configurable system components, the CritSpace Manager is used to load and unload Workspaces.

A workspace is the top-level unit of work created by users and provides the display context for rendering and interacting with panels. The base panel object provided by the CritSpace framework communicates basic information about its current state using the repository proxy. The base panel object is instantiated only through customized extensions. These extensions may communicate with content repositories that provide access to domain specific data collections.

4.1.2 Creating a Custom Deployment

CritSpace is designed to be a framework that can be tightly integrated with an existing digital library rather than a stand-alone application. The result is that CritSpace requires a custom deployment rather than simple installation. To make this process more concrete, this section describes the steps required to create a custom deployment. This process includes setting up the initial application context and loading custom panel implementations. While implementing domain specific functionality may be as complex and difficult as needed to support any given domain, this complexity is encapsulated in the implementation of panel extensions (discussed in the next section). The deployment process itself is designed to be as straightforward as possible.

```
<!DOCTYPE HTML>
<html>
  <head>
    <title>CritSpaceTC Demo</title>
    <meta http-equiv="Content-Type" content="text/html; charset=utf-8" />

    <!-- Boilerplate: Include standard CSS and JS files -->
    <link type="text/css" rel="stylesheet"
          href="/assets/reset-fonts-grids.css"/>
    <script type="text/javascript"
            src="/assets/yuiloader-dom-event.js"></script>
        <!-- ... other standard includes omitted ... -->

    <!-- Include scripts & styles for this deployment -->
    <link type="text/css" rel="stylesheet" href="assets/critspace_tc.css"/>
    <script type="text/javascript" src="scripts/critspace_tc.js"></script>
  </head>

  <body class="csdl-skin csdl-skin-crit yui-skin-sam">
    <!--  Container for CritSpace application menus and components -->
(1)     <div id="chrome">
      <!--  Menu Bar -->
      <div id="menubar">
        <span class="title"></span>
        <div id="mymenubar" class="yuimenubar"></div>
      </div>

      <div id="tabpanels"></div>
    </div>

    <!--  Container to load and display workspace content -->
(2)     <div id="workspace">
      <div id="controls"></div>
    </div>
  </body>
</html>
```

Fig. 8. The base HTML page used by the CritSpaceTC demo

The CritSpace application consists of an HTML file that will serve as the base of the application and a JavaScript file that initializes the main system variables. The base HTML page used by the CritSpaceTC demo is shown in Figure 8. The document body provides two top-level elements: 1) the “Chrome” used to display application-level interface components and 2) the workspace on which digital objects are displayed. There is some flexibility in what content is included in these components. For example, in the CritSpaceTC chrome, there is a “tabpanel” element. This is part of an (as yet unfinished) extension designed to support a general purpose Rich Text Editor that will be displayed as a set of tabs at the bottom of the application to let users write extended notes while working in with CritSpace. The specific functionality of this component will be implemented in a set of dynamically loaded script and CSS files. Custom deployments can use a similar approach to introduce other components.

The CritSpaceTC deployment gives CritSpace the full display area of the web-page to work with. In most scenarios, when manipulating visual objects in a workspace, the more available space the better. This configuration is not required, however; the system allows a deployer wide latitude to customize the Web environment in which CritSpace will be deployed. The application will use the “workspace” DIV element wherever or however it happens to be configured. Deployments are free to give CritSpace only a portion of the HTML page for its display area.

The real glue for the deployment comes from the bootstrap script file. In order to demonstrate how this deployment script can be used to customize the behavior of the CritSpace framework, Figure 9 shows the main initialization portion of the CritSpaceTC

script (largely boiler plate)[3]. This is all the code that is required to create a minimally functional CritSpace deployment. This deployment will allow users to select Panels to be added to the workspace from the menu bar, but provides no custom behavior beyond that implemented by the Panel extensions.

[3] Code snippets have been edited, shortened and organized for clarity. Production code would be similar, but provide better error handling and layout.

```
var CritSpace;

// Configuration details for panels extensions to be loaded
var panelDefns = [

      { type : "org.idch.afed.FacsimileViewer",
        js   : ["/js/IDCH/afed/FacsimileViewer.js"],
        menu : "Facsimile Viewer",
        desc : "A filmstrip viewer for working with digital facsimiles.",
        depends : ["org.idch.afed.PageDisplayPanel",
                   "org.idch.images.Filmstrip"],
        modules : ["afed"]
      },

      // details for other panels omitted
];

// configure source code dependencies
YAHOO.util.Event.addListener(window, "load", function() {
    // Set configuration paramters
    // ... details omitted ...

    // .. tell CritSpace where to find the web API for CritSpace
    $P("idch.critspace.urls.ws",     "/CritSpace/workspaces");
    $P("idch.critspace.urls.panels", "/CritSpace/panels");
    // ... details omitted ...

    // .. load module dependencies (transitive dependency management)
    var modules = ["critspace", "afed", "images-scroller",
                   "nt-tce", "tzivi", "tzivi-images", "tzivi-markers",
                   "selector", "animation"];

    IDCH.load(modules, true, config);

});

// Once all the dependencies have been loaded onto the page, this method
// configures the base CritSpace deployment. This loads the initial Panel
// extensions along with the repository to be used to access the server.
//
// Once CritSpace is loaded, this will invoke the 'main' function
// to perform the deployment specific initialization.
function config() {
    CritSpace = IDCH.critspace.CritSpace;
    var repo = new IDCH.vprops.BasicRepository(),
        config = { panels : panelDefns,
                   vpropsRepository : repo };

    CritSpace.init(config,  {
              success : main,
              failure : handleFailure
          });
}
```

Fig. 9. The initialization portion of the CritSpaceTC script

While the minimal deployment may be adequate in many scenarios, at other times deployers will want to provide a more customized environment. This might involve adding key listeners to provide users with shortcuts to add panels (e.g. ctrl-n to add a new note). Since CritSpaceTC is designed to support a specific user study, I've added customization support to initialize the workspace for the two different study groups. Figure 10 shows the code segments that implement this behavior.

With this minimal bootstrapping code, CritSpaceTC is set up and connected to a digital facsimile server. While additional tweaking and user-interface sugar may be desired, this is a fully functional deployment. The real power of CritSpace, however, comes not from the scaffolding to deploy it as an interface to a specific digital collection, but in the features that it provides for describing domain specific interactions with individual data objects and library services and the facilities it supplies for expressing relationships between those objects. The following sections address these aspects in more detail.

```
// The main function (called once the environment has been configured)
// attaches an event listener that initializes new workspaces.
function main() {
    CritSpace.on("workspaceloaded", function(ws) {
        ws.on("ready", function() { initWorkspace(ws); });
    });
}

// If a new workspace is created, this calls the appropriate initialization
// routine for study group 1 or 2 based on the workspace's name.
function initWorkspace(ws) {
    var size = ws.listPanels().length,
        name = ws.getName();

    m_workspace = ws;
    if (size == 0) {            // this is a new workspace
        if (name.startsWith("Group1/"))         initGroup1();
        else if (name.startsWith("Group2/"))    initGroup2();
    }
}

// The workspace initialization method. Specifies the facsimile to use in
// the study and loads it and the BaseText panel into the workspace.
function initGroup1() {
    var id = 7736,
        resolver = new APIResolver($P('idch.afed.facsim.image.servlet')),
        bt_cfg = { /* Base text configuration details omitted */ };

    // First, we configure and load the facsimile)
    IDCH.afed.Facsimile.get(id, function(facs, msg) {

        // Once that's ready, we load a FacsimileViewer into the workspace,
        // passing in the configuration details for this Panel Extension
        m_workspace.createPanel(IDCH.afed.FACSIMILE_VIEWER, {
            position : { y : 75, x : 50 },
            size : { width  : 800, height : 166 },
            props : { facsimileId : id },
            imageProperties : { width : 152, height  : 146 },
            resolver : resolver,
            format : "small"
        });
    });

    // load the base text panel
    m_workspace.createPanel(IDCH.nt.BASE_TEXT_PANEL, bt_cfg,
        function(panel) {
            // once loaded, set the appropriate text to display
            panel.on("ready", function() {
                panel.setChapter("Luke", 20);
            });
        });
}

function initGroup2() {
    // ... follows the same approach as initGroup1 ...
}
```

Fig. 10. Implementation of custom behavior to support the needs of the user study application

4.2 Panels

In CritSpace, users interact with content (short text notes, images, digital facsimiles, etc.) and access library services (e.g., a faceted search interface) through panels. These panels provide:

- the canvas on which extensions display content
- a general model for visual interaction
- a mechanism for simple data persistence

Much like other spatial hypertext systems, panels are graphical objects that the user works with by manipulating visual properties, such as its placement in a two-and-a-half dimensional workspace, background color, font properties, and border style. In CritSpace, however, a "panel" is an abstract entity that provides an extensible model of visual properties (discussed later in this chapter) and supports an API for transparent, server-side data persistence so that panels and their state can be saved to and restored from the remote CritSpace server—a non-trivial task in a web-application. This represents the core spatial hypertext functionality of CritSpace. The panel implementation itself, however, does not specify what content is displayed or how users should interact with that content. It merely serves as a visual canvas object that must be extended to support specific types of content and user interactions. This approach separates the spatial hypertext features that support lightweight information seeking and structuring tasks from panel "behaviors" that enable the use of that information within

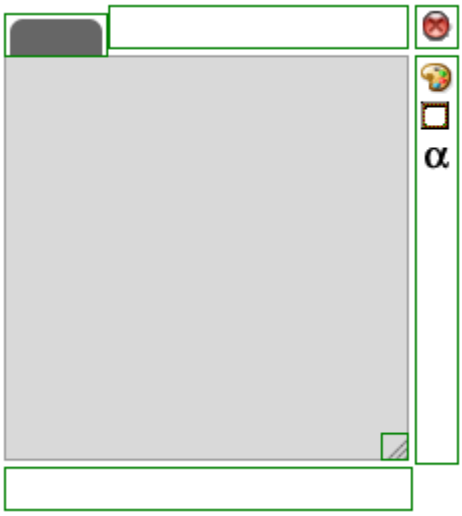

Fig. 11. A generic Panel object. Control areas are shown surrounded by green boxes for clarity.

the spatial hypertext application—uses that will often require domain-specific implementation.

Panel behaviors are implemented through pluggable extensions. These extensions use the display area (an HTML DIV element) provided by the panel to render domain-specific content and the UI controls that allow users to interact with that content. In addition to the main content area, extensions may also take advantage of control areas located around the border of the panel object. Figure 11 shows the anatomy of a panel and its control areas. Control areas include system provided functionality (such as the tab at the upper left-hand corner and the resize handle in the lower right-hand corner that support dragging and resizing of panels and the visual properties toolbar on the right-hand side) as well as regions for use by panel extensions.

I have created a number of basic and domain-specific panels. The basic panels (in the sense that they support use-cases common to many different domains) include a

text/note panel that allows users to enter free-form text and image panels for simple (e.g. jpeg) images and for tiled, zoomable images. A filmstrip panel displays large sets of images and dynamically loads and unloads images as they become visible in order to improve performance.

To support our test domain of textual criticism, I have implemented extensions to these panels that are designed specifically for working with digital facsimiles. The FacsimileViewer extends the basic filmstrip viewer and integrates with the AFED tools we have developed separately to represent digital facsimiles [Audenaert and Furuta 2008]. Rather than storing information about the digital facsimile in the CritSpace system, the FacsimileViewer uses the AFED web API to retrieve information about the digital object to display and stores the facsimile id using the panel's named properties. A PageDisplayPanel extends the basic image panel to provide access to high-resolution, zoomable page images from the digital facsimile.

Base panel objects are created and managed by the CritSpace framework and extensions are plugged into these objects at run-time, dynamically adding methods and properties to base panel objects. Extensions can be implemented to depend on other extensions in order to promote flexible code reuse. For example, the FacsimileViewer adds the plugins for an AbstractFacsimile object and for a Filmstrip panel before adding in its own methods. In this case, the FacsimileViewer extension simply provides the programmatic glue that translates between the AFED domain model of a facsimile and the filmstrip interface widget. The result is a sophisticated interface component implemented in fewer than 150 lines of code.

4.2.1 Extending Panels

A key motivating factor for CritSpace is the need to supplement the strengths of spatial hypertext systems for dealing with ambiguous and partially formed knowledge with explicit support for the data cultural heritage (and other) scholars use and tasks they perform. CritSpace addresses this need by distinguishing between the visual properties of content object and the behaviors of those objects. Behaviors are implemented as extensions of the core panel objects that are instantiated and maintained by the framework.

A panel extension (henceforth, an Extension) is implemented as a JavaScript function[4]. This function takes two parameters, a base panel object that has been created by the framework and a set of configuration parameters. An Extension extends a panel in two primary ways. First, it can create private functions that write HTML content to the panel's display canvas (a DIV element). These functions can be attached as event handlers to respond to system level events (e.g., when a user selects a panel or a visual property changes), domain level events (e.g., when content is loaded from a server), or HTML DOM event (e.g., when a user clicks on a button that the panel added to the display). Second, it can attach functions or objects to the base panel object in order to expose functionality and data to other components of the system.

---

[4] As a functional programming language, JavaScript functions create a closure that, when the function is invoked, instantiate their own internal state. Objects created in this closure can be accessed only by other objects inside the closure unless they are made public, for example, by returning an object from the function call or attaching it to another object. Functions themselves can be treated like any other object in the program—they can be assigned to variables or attached as members of other objects.

Extensions can be arbitrarily complex and can rely on external libraries. Extensions can also build on other Extensions. Since an Extension is simply a function, it may simply invoke another Extension, passing in the base panel object and configuration details that it received (possibly after modifying those configuration details as needed)[5].

Consider, for example, the FilmstripPanel Extension. This relies on a separately developed image scroller widget to display very large sets of images. The scroller takes a list of image identifiers, and an object that tells it how to translate those identifiers into URLs. It then loads the images into a scrollable DIV. Since loading large numbers of images onto a web page will quickly degrade performance, the scroller dynamically loads and unloads images as they become visible (plus a buffer of some images that are just out of view). The scroller lets users select images and drag images in order to re-order the list or to remove them by dragging images outside of the scroll area. Each of these user interactions fires an event. The default behaviors can be customized by setting configuration parameters or by listening to these events and taking any required actions.

The FilmstripPanel Extension provides the glue that connects this independent interface component into the CritSpace framework. When the Extension is first invoked, it creates an instance of the image scroller and renders it using the canvas provided the base panel. It then attaches functions that listen for user interactions with the scroller.

---

[5] This is not presented as the best model for extending functionality, but rather one possible implementation of this core concept. Plans for future improvements include a more robust implementation of support for extending base functionality and mixing in specific functionality.

Among the things that it listens for are drag events. When an image is dragged outside of the panel area onto the workspace, the filmstrip creates a new ImagePanel.

One noteworthy feature of this approach to extending panels is that individual Extensions can (should) be very specific. The FilmstripPanel, for example, displays a large set of images. It says nothing about how to obtain those images. To be used, it needs either to be extended further or else used by another panel (for example, an image search panel might use a FilmstripPanel to display its results) in order to obtain images to display.

In our test case implementation, CritSpaceTC, the FacsimileViewer panel provides the functionality for retrieving digital facsimiles of manuscripts from a digital collection and relies on the FilmstripPanel to display page images. The `FacsimileViewer` overrides some of the basic functionality of the filmstrip such as preventing users from re-arranging or deleting page images and loading a PageDisplayPanel instead of an ImagePanel when the user drags an image onto the workspace. The pluggable extension model promotes modular development patterns that encourage code reuse and separation of concerns. It also allows panel extensions to easily build on existing components as demonstrated by our use of the image scroller widget and the AFED system.

Creating panel extensions is only part of the solution. Once created, the system needs to be able to discover these extensions, load panels into the workspace as needed, and restore them when a user loads a previously used workspace. To achieve this, CritSpace provides a simple extension registration service. The developer in charge of

creating a CritSpace deployment, identifies the Extensions to be used in the deployment and registers them as part of the initial page loading script.

Panels are instantiated in the workspace in one of three ways: users can select a panel type from the CritSpace menu, an Extension can be designed to create a new panel objects, or through an *ad hoc* bootstrapping script. In all cases, panels are created by calling the API of the Workspace object and passing the type of a registered panel and any custom configuration parameters. The easiest way to get a panel into the workspace is through CritSpace's build-in menus. The registration systems allows panel Extensions to be added to the application menu when they are registered. This is envisioned to provide access to search and browse panel Extensions that will be used to generate other panels in the workspace. Once a panel has been loaded, it can use the Workspace API to generate new panel objects in response to user interactions (e.g., dragging an image from the FacsimileViewer onto the workspace).

Alternatively, panels can be created programmatically by the bootstrap scripts that deploy the system. This could be done, for example, by attaching event listeners that respond to keystrokes or mouse events by creating a panel. For example, a deployment might create short-cut keys to allow quick entry of bibliographical notes or respond to double-click events by creating a new panel for simple text entry. In some cases, it might be appropriate to load workspaces with an initial context, such as a panel to search a particular digital library's online collection. The CritSpaceTC system initializes all new workspaces by creating an initial FacsimileViewer and BaseTextPanel.

4.2.2 Managing Persistent Data: Named Properties

In order to display content and to allow users to access that content over multiple sessions, panel extensions need to work with two different sets of data: (1) the source data representing the content that is displayed to the user and (2) state or session data that represents the current state of the display. Typically, source data will be taken from an external digital library source. In our test domain, this includes things like the images and metadata for a digital facsimile, the base-text and any edited forms of that base text that reflect the wording in a specific manuscript. This source data is outside the scope of CritSpace. It is consumed and manipulated by panel extensions, but the core framework provides no support for data storage or access.

To manage access over multiple sessions, however, CritSpace needs to represent the current state of the data that is being displayed in a workspace. Session data includes information about what content is being displayed, for example, the identity of the digital facsimile and the specific image that has been selected. It may also need information about how that content is being displayed such as the zoom level and position within the panel of an image. The specifics of this session data will be heavily dependent on the implementation of individual panel extensions.

CritSpace supports persistent session state through named properties. Named properties allow panel extensions to store state information using a simple, *ad hoc* key/value map. The core panel implementation provides methods to set, retrieve, and list the named properties that are associated with each panel in the workspace. CritSpace automatically persists these properties to the server-side storage layer, thereby relieving

panel extension implementers of the need to handle the myriad of details required to persist data in a web-based environment. Furthermore, future extensions that support tracking user history can be automatically integrated with little or no change required to the extension itself.

Named properties are similar to the "slots" of the early VIKI system [Marshall et al. 1994] with one significant difference: named properties are not intended for direct access by end users. Direct manipulation of metadata values requires a degree of explicit representation of knowledge that seems at odds with the informality that lies at the heart of spatial hypertext systems. While slots can support a very general model for incremental formalization, the practical value of providing users with direct access to such an open-ended formalization mechanism is unclear. Instead, I chose to mediate access to this feature through the programmatic control of the panel extensions. Thus, where fully general support for user supplied metadata is needed, it is trivial to implement an extension that provides users with direct control over "slots" in a VIKI-like interface. Otherwise custom extensions can simplify access to relevant named metadata properties or provide explicit support for common information organization practices that are common to a domain such as tagging or note-cards formats.

4.3 Links

Links are a key feature in any hypertext system. They typically support navigation between information nodes (e.g., pages on the web), semantic relationships, or, as in the case of the semantic web, both. Early interactive visual interfaces envisioned

map-based hypermedia systems as a mechanism to allow users to express formal relationships between information nodes, this research found that users were unwilling or unable to create those explicit relationships [Marshall and Rogers 1992]. Subsequent spatial hypertext systems eschewed explicitly specified links, relying instead on algorithmic processing to recognize implicit linkages that are visually apparent to human readers but not formally expressed in terms of the system's representation of the space [Shipman et al. 1995]. Later, VKB re-introduced a navigational linking model to allow users to move between different regions of the hypertext (or even between different versions of the hypertext's history) [Shipman et al. 2001a], but spatial hypertext systems continue to focus primarily on implicit rather than explicit linkages.

CritSpace enables explicit linking between panels, but unlike links commonly found hypertext systems that establish navigational or semantic relationships between content nodes, links in CritSpace specify behavioral relationships. Users (and the developers who implement extensions) work in contexts where different chunks of information are related. Sometimes those relationships are implicit and well supported by the spatial and visual organization strategies. At other times, however, those relationships affect the actions that can be performed on some content. Those relationships may also mean that actions taken on one object affect another object.

For example, when a user drags a page image from the FacsimileViewer onto the workspace, the FacsimileViewer creates a new PageDisplayPanel. These two panels are related in the sense that they both provide a view onto the same underlying domain object, a particular digital facsimile stored in a particular collection. This "linking by

reference" in which two panels provide different views onto the same underlying object requires no explicit support from CritSpace; it is implicit in the implementation of these two panels. Both of these panels also maintain internal state about the page that is currently selected or displayed. By default, these panels are linked so that changes to one panel (selecting a new image in the FacsimileViewer, for example) is reflected in the other panel. This allows the PageDisplayPanel to serve as a detailed view pane that corresponds to the overview provided by the FacsimileViewer.

This creates a linkage between the behavioral aspects of the panels that is not specifically a property of the underlying data model. Instead, these linkages represent interaction design decisions made by the implementer of the panel extension. Unlike the implicit links between panels and the data that they display, these behavioral links are contingent; a user may choose to break the connection, for example, by decoupling the PageDisplayPanel from the FacsimileViewer in order to compare multiple pages side by side.

Unlike the connections driven by the underlying domain model of the source data, this sort of transient, interaction driven link requires explicit support from the system, which CritSpace provides via the PanelLink mechanism. Much like panels, PanelLinks in CritSpace define the existence of a relationship between two or more panels and persist that relationship between sessions. They do not, however, specify how that link affects the behavior of the linked panels.

PanelLinks can also be used to alter the behavior of existing panels. The BaseText panel implemented for CritSpace provides a thin proof-of-concept UI widget

designed to support textual editing. It has been designed to work with the more mature AFED module, but was developed separately. The implementation of this panel extension includes a PanelLink that connects a BaseText panel to be connected to a PageDisplayPanel. This link modifies the behavior of the PageDisplayPanel, attaching a reference to the zoomable image viewer to the base panel object. This image viewer allows for multiple, synchronized overlays to be displayed in the same frame.

The BaseText link takes advantage of this to add a new display layer to the image. This layer allows a user to add verse markers (in this case, the example domain is working with Greek New Testament manuscripts rather than general purpose textual criticism) by pressing control and clicking on the appropriate location in the image. These verse markers are initialized based on the text in the base text panel (for example, if the base text is Luke 2, the first marker will be for Luke 2:1 and subsequent markers will increment sequentially). These marks are linked to the base text; clicking on the verse marker causes the appropriate verse in the base text panel to be highlighted and scrolled into view. The user can break this link, thereby removing the image overlay and stopping the synchronization behavior. A complete implementation of this component would rely on a server-side text-to-image linking tool to store the domain-specific relationship between the image and the base text and would load the appropriate data into the overlays based on data from the server.

This example implementation illustrates how the PanelLink mechanism supports general purpose linkages between panels. These linkages are not merely navigational or semantic (although, in this case they are both), but serve primarily as a tool to enable

developers to provide support domain-specific user interaction. These interactions allow the system to take advantage of data-driven structure and out-of-band domain knowledge to support scholarly needs for analytical tasks while minimizing the cognitive overhead imposed by working with the application.

4.4 Visual Properties

A central component of any spatial hypermedia system is the facilities it provides for manipulating the visual properties of the objects users interact with. These properties typically include the position on a canvas, background color, border style, and font properties. While most systems provide a fixed set of visual properties, the framework-based approach adopted by CritSpace makes it, in the best case, highly desirable that externally developed extensions that build on the default set of visual properties. For more innovative uses, these extensions are not a matter of convenience, but are required to implement certain functionality. A comprehensive set of visual properties that covers all possible scenarios is not possible. Indeed, I have found that even in relatively simple scenarios such as implementing support for zoomable images, requires careful consideration of how visual properties such as width and height behave.

A fixed set of visual properties, even if large and well designed, will inevitably limit the scope of interactions that can be implemented with CritSpace. To avoid this, while still providing system-managed support for spatial hypertext, CritSpace includes a module (called vprops) that defines a high-level model of visual properties while allowing implementations to selectively include different visual properties or create new

properties. Since CritSpace is a web-based application, this module is designed to work easily with CSS-based properties, but can be extended as needed for specific use cases.

The chief design goal of vprops is to provide strong support for persistent visual properties that meet the majority of needs and can be extended to handle new needs as they may arise. For example, recent implementations of VKB have introduced shadows that are used to express system information derived from the spatial parser or other computational analysis of a workspace [Bae, et al. 2010]. The system can express information by changing the color and opacity of these shadows. Alternatively, the current CritSpace implementation represents background color using an RGB color pallet. This is motivated on the idea that colors express named concepts—yellow messages are informational, red messages are warnings, for example. An alternative approach might be to use the background color to represent two dimensions of information, one expressed through the hue of the background, the other expressed through its saturation.

Vprops was designed so that alternative strategies for implementing the visual properties of a spatial hypertext system could be achieved entirely in the interface layer, without requiring changes to the underlying data model used to represent those properties. This is particularly important with CritSpace since, as a web-based system, adding a new type of visual property would require changes at the persistence layer, the web API layer and the AJAX communications between the client application and the web-server. vprops allows the introduction of new properties such as those described

above by making comparatively minimal changes to the JavaScript code that powers the client interface.

In designing the visual properties module, I have identified the following major objectives:

Simple data persistence vprops needs to provide a simple, transparent mechanism for data persistence that abstracts the process of AJAX based communication with the server as well as interactions with the underlying database. This allows developers and designers to work with visual properties as if they were plain-old objects, greatly simplifying development tasks. To improve performance, this persistence uses optimistic error handling strategies. Properties report and attempt to recover from errors if they occur, but act as if commands succeed when they are executed rather than waiting for the round trip to the server to complete.

Grouping and subgrouping Visual properties often represent natural groups. For example, borders include properties for color, style, and width. Positional properties include x, y, and z coordinates. vprops should supports hierarchical organization of properties for convenience.

Default and preview values vprops should support default values for all properties. Default values go beyond providing sensible initial values (though they do that), defaults also provide a mechanism for the system to gently manipulate the visual display without overriding a user's manual settings. Properties should also support the notion of a transient preview value. Interface controls frequently need to reflect potential changes to a user (for example, showing a new background color) before the user

decides to make that change. Others may need to temporarily modify an object to call attention to it, for example, changing a border to red to indicate that the panel is an eligible link target. Since property values are transparently persisted, a separate preview mechanism allows these transient display changes without introducing un-necessary network traffic or invoking other system components such as a history mechanism[6] that might be listening for changes to visual properties.

Constraints on values A visual property definition should be able to specify constraints on the values that are valid for that property. These constraints may be fixed as part of the property definition (for example, color values must represent valid hexadecimal color codes) or may be driven programmatically based on the state of content displayed in a panel (a pane to display an image should maintain the aspect ratio of the image). In addition to identifying valid and invalid values, properties should be able to map invalid values to valid ones where such mappings are possible. Notably, this feature depends on programmatically controlled behavior and requires that the persistent state of visual properties can be matched with the runtime implementation of that behavior.

[6] History management is not part of the current implementation but is envisioned for future work.

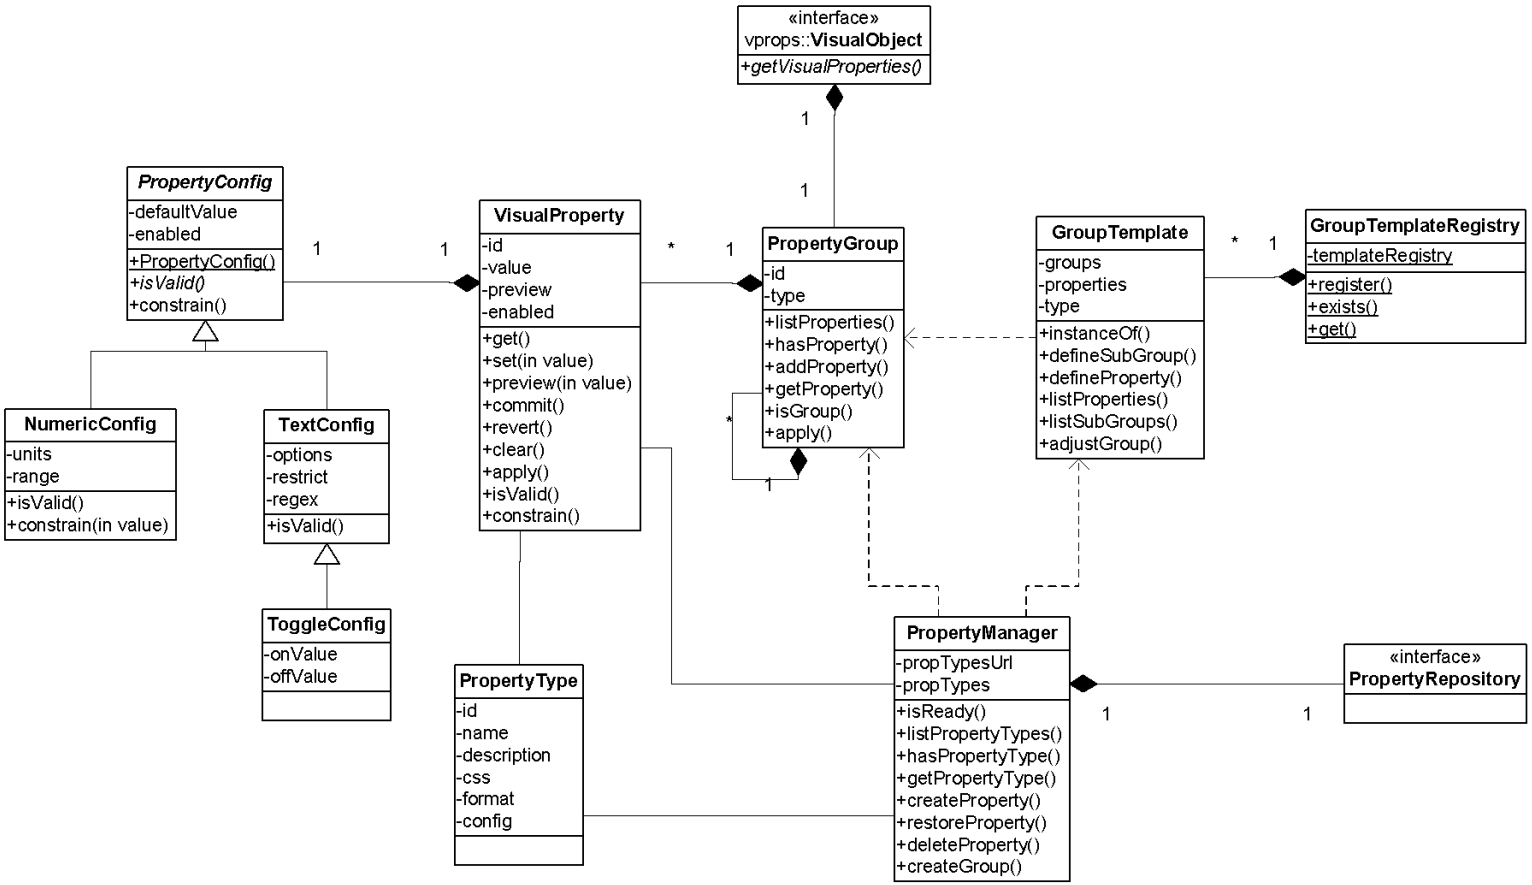

Fig. 12. UML diagram of the visual properties module

Figure 12 shows a UML diagram of the vprops module architecture. This system has no dependencies on the rest of the CritSpace framework and so could (conceivably) be deployed separately to model visual properties in other applications, if warranted. The system is divided into a server-side component that manages the persistent storage of visual properties between user sessions and a client side component implemented entirely in JavaScript.

The server provides a simple HTTP API and the client communicates with the server exclusively through a Repository object. Encapsulating all client-server communication in a single class allows alternative back-end implementations to be plugged in. These implementations could take advantage of advanced client-side caching provided by modern browsers, interface with alternative implementations of the server-side component, or provide a stripped down version that doesn't perform remote API calls for use in unit testing.

Dynamic loading of custom properties is accomplished through the use of a properties registry. All visual properties implement a common data model and supply a run-time implementation that determines how to interpret and apply this data model to a panel. Both default and custom property implementations and loaded into a visual properties registry by the CritSpace deployment. When a panel is loaded into the application, CritSpace attempts to load any defined properties for that panel from the registry. Since updates to the deployment may mean that a visual property that was available can no-longer be found, missing visual properties simply are not rendered.

CHAPTER V

EVALUATION

CritSpace is proposed as a platform for building domain-specific, interactive visual interfaces that support scholarly practice. To evaluate this claim, I conducted a focused case study that walked through an abbreviated design, deployment and evaluation cycle. This study was conducted in the context of an existing collaboration with the H. Milton Haggard Center for the Study of New Testament Manuscripts (CNTTS). While the current work is purely research oriented and will not be deployed in direct support of CNTTS's daily research operations, this work was conducted with a view to evaluate technologies that may be deployed in production in the future.

CNTTS is a research center at New Orleans Baptist Theological Seminary devoted to the study of the Greek New Testament text. Founded in 1998, it has become one of the premier settings worldwide for studying the Greek New Testament text. CNTTS has created the world's first comprehensive, searchable electronic database of variants in the entire Greek New Testament. The *HCNTTS Critical Apparatus* includes variants of all types (including minor differences such as orthographic shifts, and nomina sacra) with a sampling of at least 50 manuscripts for each book of the New Testament.

At CNTTS, a staff of seven graduate students and a visiting scholar work under the direction of the center director on a variety of tasks focused on the detailed examination of Greek New Testament manuscripts. Over the past decade, the primary focus of the center has been on the comparison of manuscripts (typically printouts of

microfilmed or digital images) to an established base text in order to produce a comprehensive catalog of differences. For the purposes of evaluating CritSpace, this center offers a large number of scholars actively engaged in work with source documents, many of which can be easily added to a prototype digital library tool for the purposes of a study.

The case study consisted of three main stages. In the first stage, I conducted week-long pre-study consisting of contextual inquiry and observation. During this time I gained an understanding of both their scholarly objectives and their work practices. Based on these observations, in the second stage I designed and implemented three new CritSpace panels that reflect their current paper-based work-practices within the CritSpace environment. This design benefited from ongoing interaction with CNTTS during the design and implementation phases. Finally, I presented this interface to 22 researchers and graduate students affiliated with CNTTS. These participants performed collations of two chapters from the book of Luke using both the traditional paper-based workflow and the CritSpaceTC system. In addition to observing each participant during the study task, I concluded by conducting a short semi-structured interview about their experience with the tool.

5.1 Pre-Study

Prior to selecting CNTTS for the case study to evaluate CritSpace, I conducted an intensive formative study consisting of in-situ observation and contextual inquiry as part of a separate collaboration. This study is helpful for providing an in-depth analysis

of a mature paper-based workflow with a view toward understanding how that workflow can be adapted to the context of the CritSpace framework.

5.1.1 Methodology

Over the course of this four day study I conducted informal interviews six graduate student researchers in the center (all Ph.D. students at the seminary), a student worker (who primarily assisted with creating digital facsimiles from microfilms and other office work), the visiting scholar and the director. During these interviews we discussed the division of labor among center staff, general and specific work practices, training procedures and other topics related to collating manuscripts and creating the apparatus. The duration and informal nature of this study allowed me time to reflect on my observations and ask follow up questions as needed. I took notes on paper, but the interviews were not otherwise recorded.

In addition to these interviews, I conducted all of my work (reviewing my notes and planning for further interviews) at the center. This allowed me to the freedom to work in the same environment as the center staff and observe their work practices unobtrusively. I also attended a class on advanced textual criticism taught by the center director and observed an "all hands" group meeting. During the time I was conducting these interviews and observations, the center was finalizing the last collations for a major update of their apparatus in preparation for a final review and release to Oak Tree Software, Inc. for publication as a module in their *Accordance* software package. This work was completed in time for the group meeting on the morning of the fourth day,

which I attended. This scheduling coincidence ensured that there was plenty of focused activity for me to observe.

In additions to these interviews and observations, I examined a detailed description of the critical apparatus and reviewed their source data files as well as the *Accordance* software package. To validate my understanding of their workflow and the domain model of the apparatus I prepared a number of rapid interface design sketches, web site workflows and informal data models that I reviewed with senior students and the center director in order to ensure that I had correctly understood their work practices and needs.

5.1.2 Observations

CNTTS relies on a paper-based collation process. A research assistant (the collator) typically begins with a printout of a manuscript taken from either their microfilm collection (now digitized) or a digitally photograph facsimile produced by a related organization. The collator compares this manuscript to a printed base text, the UBS4/NA27 critical edition of the Greek NT. This base text is printed double spaced, with lined space in the right-hand margins for notes or to record manuscript features. The first page of the base-text printout provides space for recording information about who performed the collation, what manuscript was used and when the collation took place.

The collator begins by marking verse breaks in the manuscript images (since verse divisions were added long after these manuscripts were created) and the line

breaks from the manuscript in the base text. These marks make it easier to find the location of a word or phrase from the base text in the manuscript and vice versa.

For the collation process itself, the collator compares the two documents word by word. Any discrepancies are noted in the interlinear space on the base-text. For variants that do not fit between the lines, the ruled space in the margin may be used for additional notes. With the exception of accent marks and punctuation, both relatively late additions to Greek writing, all variants are noted. Unique scribal abbreviations or short-hand and unusual glyphs are reproduced by hand on the base-text copy. The collation is typically performed sitting at a desk. The stacks of unbound paper for the manuscript print-out and the base text are placed side by side on the desk. As the collation proceeds, the collator turns the completed pages face down on the desk, creating a line of four stacks of paper. The inner two stacks are the "todo" pile and the outer two are the completed work. While this procedure is the most common, a few of the collators use computer-based images of the manuscript in favor of printouts (while still noting variants on a printed copy of the base text).

Each manuscript is collated by two different graduate assistants working independently. Once both collations have been completed, one of the collators (or another member of the center) compares the two collations for any discrepancies. These discrepancies are checked against the source manuscript and the corrections noted on one copy of the base text in a different color of ink. Once this reconciliation process has been completed, the corrected base text is given to the project manager (PM) who updates the apparatus.

5.1.3 Discussion

Despite its apparent simplicity, collation is not a task that can be reduced to mere computation (by either humans or machines). Collators are required to make informed judgments about the readings of difficult passages based on their knowledge of the language, scribal practices in general, the tendencies of the specific scribe who copied a manuscript, and their prior experience. The paper-based methodology that CNTTS uses has evolved to facilitate this process over the course of a decade's worth of experience working with NT manuscripts. The optimizations are subtle: lines in the margin of the base text, starting off by marking the verses and line breaks on the respective print-outs, the uniform use of the four-stacks-of-paper method for managing the data during the collation, the use of a different color of ink to record the reconciled transcription. Many of these seem obvious in retrospect, but they are innovations introduced through hard won experience.

A few points are worth highlighting. The collation task is one that requires careful attention to the spatial arrangement of the materials being used. The organization of printouts and the use of visual cues in the forms of verse markings and line breaks helps participants to keep track of where they are in the collation process. This spatial arrangement is not universal; the collators who preferred to work from digital images arranged their workspace differently from those who work exclusively with paper, but they too had a specific spatial arrangement of material that they used consistently in order to facilitate their work. CNTTS had previously experimented with software support for their collation project, but found that, among other issues, the context

switching that was required between the (at the time) paper-based manuscript images and the text on the screen lowered accuracy rates. The ability to arrange objects spatially in the CritSpace system may alleviate some of these previously encountered difficulties.

Complicating this, however, is the fact that unlike more exploratory research questions, the collation of variants is a highly structured task. Spatial arrangement of objects is a background task, rather than a foreground task as it is in more traditional knowledge organization and structuring applications of interactive visual interface. Consequently, this study will depend more on CritSpace's ability to support linked interactions between and sustained simultaneous use of multiple display objects. In addition, to the extent that the study participants do rely on the spatial features of the system within the context of this highly focused work, these features should prove to be even more promising to support more open-ended use cases.

Finally, to take a step back from the details of the collation process, a couple of features of this work make the textual criticism (TC) well suited to this initial evaluation work. As noted in Chapter III, variants of this work are a common reason that scholars need access to source documents instead of modern editions. The nature of this work allows for the creation of short-duration (45 minutes) tasks that both realistically represent scholarly work practice and yield meaningfully measurable results. And, owing to the fact that TC on this scale requires many people to be involved in the collation process, this domain affords the opportunity to study several people with varying levels of expertise working on an identical problem. While further evaluation

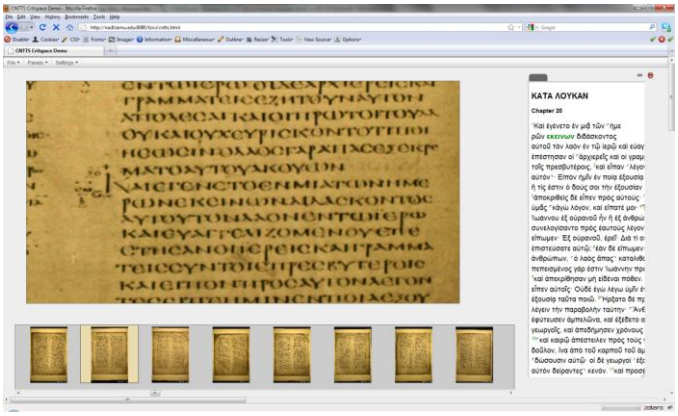

Fig. 13. Designing CritSpaceTC to mimic observed paper-based collation practices.

under more open-ended cases will be needed, TC provides an exceptional use case to start with.

5.2 Design of CritSpace Extensions

Building on the results of these initial observations, I designed a custom deployment of CritSpace deployment called CritSpaceTC. As shown in Figure 13, CritSpaceTC is intended to replicate the paper-based workflow of CNTTS as closely as possible. It provides the user with a side-by-side view of the manuscript and base text. Users will then edit the base text in order to mark the variations that they find in the manuscript. Many of the features of the specific panel implementations have been introduced in the previous chapter as examples of CritSpace's extensibility. In this section I discuss them within the context of supporting New Testament TC.

5.2.1 The Panel Extensions

CritSpaceTC introduces three custom panels: a FacsimileViewer, a PageDisplay and a BaseText for the collation. These panels display structured data drawn from multiple sources (facsimile documents, the base-text transcription, and the critical apparatus) and introduce user interface features to manipulate that content. The panels rely heavily on event passing and an extensible panel API to support communication between panels.

The FacsimileViewer makes use of the AFED module [Audenaert and Furuta 2008], a digital facsimile management system I have developed separately. This panel loads a facsimile into the workspace and displays the page images in a filmstrip-like viewer. Clicking on an image selects that page and the panel displays a brief description of the page in the area just below the panel.

The PageDisplay panel provides allows detailed inspection of high-resolution images using a tiled, zoomable image display. It too draws its source data from the underlying AFED representation of the facsimile. This allows users to sequentially navigate a document within the page display viewer. The PageDisplay panel supports adding new information overlays onto the image being displayed. In CritSpaceTC, I have used this to display verse markings on the manuscript images.

Finally, the BaseText panel allows users to mark the Greek New Testament base text [SBLGNT 2010] in order to indicate variant readings. This panel provides the collator with a heavily customized rich text editor. This editor allows the user to mark readings by highlighting a section of the base text and opening the variant editor

window. This window lets the collator enter the observed reading or else mark the reading as omitted by leaving the text blank. The panel also allows the user mark new lines by ctrl-clicking on the position in the text where a line break occurs in the manuscript.

Since the emphasis is on working with manuscript rather than organizing information, the visual properties that are exposed to users are minimal. Users can adjust the spatial position of the panels and their height and width. Extensive use is made of the constraints that can be applied to visual properties (and event listeners to respond to changes in those properties) in order to ensure that the resizing behavior of the panels makes sense given the content those panels display.

5.2.2 The Links

Collation work, in contrast with information triage, is not a process of organizing semi-independent chunks of information but rather of focused work with closely inter-related digital objects and multiple views onto a single object. CritSpaceTC is facilitates that through extensive use of CritSpace's linking facilities.

The FacsimileViewer is used to create the PageDisplayPanel (by dragging a thumbnail image onto the workspace). This process also sets up a behavior link between the two panels so that changes to the page in either the FacsimileViewer or the PageDisplayPanel will update the corresponding panel. This implements a focus+context visualization pattern [Card et al. 1999] that allows the collator to gain some perspective on his place in the facsimile and to quickly navigate to near-by pages. This context

stands in for the stacks of previously seen pages used when working with facsimile printouts.

The BaseTextPanel can also be linked to the PageDisplayPanel. This linkage is used by the BaseTextPanel to add an overlay to the page image display that lets the user add verse markers. This overlay is created and managed by the link object and is synchronized with the base text. Verse markers are added sequentially based on the text display in the BaseTextPanel, and clicking on a verse marker on the page image causes that verse to be highlighted in base text.

These linkages rely heavily on the features for prototype inheritance of behavior. For example, both the PageDisplayPanel and the FacsimileViewer extend both an abstract FacsimilePanel that defines behavior for displaying a facsimile document with multiple pages one of which is currently selected page. The base facsimile panel also defines events to notify listeners when the currently selected page changes. The synchronization link that facilitates the focus+context pattern relies on this higher-level interface rather than the details of the concrete panel implementations, thereby allowing multiple views to be synchronized. Similarly, the link between the BaseText and the PageDisplayPanel relies on a combination of a zoomable image display interface and the page selection events of the base facsimile panel in order to render and synchronize the verse maker overlay.

5.2.3 The Workflow

These panels and links have been designed to support a basic workflow that is modeled after the paper-based process currently employed. While individual use will vary, the following is presented as a basic outline of the collation process using CritSpaceTC.

The workspace starts with the FacsimileViewer and BaseTextPanel displayed. In production, these would be created by using the system menus to select a facsimile to collate and by creating a new BaseTextPanel for use in this collation. The collator then proceeds to create new PageDisplay panel by dragging an image from the FacsimileViewer onto the workspace. Once he has located the image of the appropriate manuscript page, he then links the base text to the page display. This process may proceed sequentially (by first marking verses and line breaks for a section of text and then proceeding to identify variants) or iteratively (marking verse, line breaks and variants as the collator works through the text word-by-word).

5.2.4 Discussion

The design process offers some initial insight into how the spatial interactions facilitate this domain and begins to suggest ways that CritSpace might be further adapted to better leverage the side-products of the drudge work of humanities scholarship.

On the one hand, the spatial malleability of CritSpace's interface creates an environment that users can tailor to suit their own working habits. This space isn't used for organizing information, but rather to create idiosyncratic layouts that fit individual

work preferences. The opportunity here goes beyond setting a few personalization parameters to *ad hoc* organization of the workspace. In practice, it can be expected that a few basic organizational strategies will be adopted for people working on any given task—in much the same way that only a few basic layouts were used by collators working with paper-based printouts. The key observation here is that the workspace gives users (and ultimately communities of users who informally share successful strategies amongst themselves) a tremendous degree of freedom in finding the shape that works for them.

There is a range of intellectual by-products that are created (and subsequently lost) as a result of the collation process. The focus of CNTTS is on the compilation of a critical apparatus, but that process also produces a detailed alignment of verses to manuscript images as well as identifying the line-breaks for individual manuscripts on the transcriptions. Over the course of creating their apparatus, CNTTS has produced this data for several hundred manuscripts all duly recorded as marks on paper stored in a half-dozen or so file cabinets. This information is inaccessible to the broader scholarly community and, for most practical purposes, to CNTTS itself. If no one were interested in such linkages, it could be easy to imagine this as no small loss. Linkages between transcriptions and their manuscript images, however, are useful in a variety of applications ranging from training data for document image analysis and OCR (see Rothfeder, et al. for a discussion of this [2006]) to interactive document readers such as that used by the Codex Sinaiticus project [2009].

For New Testament manuscripts, these linkages are important enough that the Institute for New Testament Textual Research (Institut für Neutestamentliche Textforschung, INTF) has created an interface used by scholars around the world to identify the starting and ending verses on each page of a manuscript. A senior research assistant at CNTTS has one of his primary responsibilities, the task of supplying this information for the manuscripts that CNTTS has access to. This work is performed separately from the collation process.

A key design feature of the CritSpace framework is its ability to draw from and, with suitable permissions, update externally developed digital repositories. While the CritSpaceTC system focused on prototyping the UI layer, this model assumes the eventual development of backend components that would include a web-based API for digital libraries (potentially housed at multiple institutions) and for a critical apparatus. The ability to integrate a third back-end service such as that provided by INTF could enable CNTTS to capture the text-to-image linking that is currently a discarded by-product of the collation process and to do so at a level of detail and quality that is too labor intensive to be a project in its own right. Thus, a relatively minor change in the CritSpaceTC interface could leverage work that is already being to save multiple person-years of effort.

5.3 Summative Evaluation

For final stage of the evaluation process, I conducted a user study of CritSpaceTC. The goal of this study was not merely to assess usability—the system is

still an early prototype application and there are many shortcomings in the area of usability—but rather to evaluate CritSpaceTC's effectiveness at mimicking and improving on the paper based workflow and to gain initial insight into the utility of the spatial manipulation of objects as a tool to support focused work within humanities disciplines.

To measure CritSpaceTC's performance relative to paper-based work, I collected both quantitative measures of performance (speed and accuracy) as well as qualitative measures such as enjoyment of the process and confidence in their collation results. The success criteria is not wholesale improvement over the paper based workflow, a goal that seems impossible to achieve given the ease of use of paper and the participants' long history of working with it. Instead, CritSpaceTC will be judged (provisionally) successful if it yields suitable performance (equivalent accuracy with some decrease in speed) and points the way toward future enhancements that can further improve both of these metrics.

With respect to the use of spatial features, the primary study conditions (performing a collation using CritSpaceTC) provides an opportunity for me to evaluate participants' use of the spatial features of the system and to ask follow up questions. These features are new to participants; they will not be explained or introduced in any level of detail; and participants are likely to focus on the collation task that with which they are familiar. Consequently, I do not expect that spatial manipulations features will play a major role in the work at hand. Instead, I hope to see examples of subtle of spatial organizations performed more or less un-consciously, and largely forgotten by the users.

That is, I hope that one of the primary contributions of CritSpaceTC will recede into the background where it provides quiet support without interrupting the scholar's real work.

5.3.1 Methodology

For the study, participants were asked to perform collations of two chapters, Luke 20 and Luke 22, both from the same Majuscule manuscript (GA02, Codex Alexandrinus). One collation was performed on paper using color printouts of the manuscript; the other was performed using the CritSpaceTC application. Participants had approximately 45 minutes to complete each task.

This work took place at the main CNTTS office space, in the same environment that CNTTS staff normally perform collations. Consistent with normal working conditions, there were numerous people present at all times during the study. Non-participants were engaged in individual and group work as well as conversations about coursework or other unrelated topics. Study participants would frequently engage in these conversations. On occasion, participants would ask others working in the center for assistance during their collation work. Since the goal of this study was to evaluate work under realistic working conditions rather than a carefully controlled laboratory, the study proctor did not intervene in these natural interactions.

The study itself consisted of a brief pilot phase with two participants on the first day, followed by 20 participants in the main study over the course of the following three days. The initial pilot phase involved two moderately experienced collators and was designed to ensure that the allotted times were sufficient, that there was an adequate

density of variants in the selected text. It also helped to identify a few areas of the code that needed minor adjustment which were made before the main study began the following day. The pilot study participants performed the collations using CritSpaceTC, but did not complete the paper based portion of the task.

The main study participants were divided into four groups (1A, 1B, 2A, 2B), roughly balanced for expected experience and skill levels. This division controlled for different relative difficulties between the two passages (group 1 collated chapter 22 on paper and chapter 20 on the computer, group 2 did the reverse) and for learning effects (sub-group A collated on paper first, sub-group B collated on the computer first). At the beginning of the study session, each participant was given a brief overview of what they would be asked to do and, after consenting to participate, filled out a short demographic questionnaire asking about their role at the center and their experience collating Greek New Testament Manuscripts. For each task, they were given 45 minutes to collate as much of the text as they were able to complete. Following each collation, they were given a follow-up questionnaire. At the end of the study, they completed a brief interview in which they were asked to compare the collation experience between paper and on the computer. The entire study typically lasted two hours.

Prior to performing the computer based task, participants were given a short (approximately 10 minutes) introduction to CritSpaceTC in which the proctor showed them the major interface components and demonstrated how to use the BaseText panel to record variants. Participants did not have any direct practice with the system prior to beginning the task, but were allowed to ask questions at any time during the study if they

needed specific reminders about how to perform some action. At the beginning of the computer-based task, the proctor set up the workspace with the base text displayed and the correct manuscript page shown in the PageDisplayViewer. In both the computer-based and the paper-based tasks, the proctor indicated the correct starting position on the manuscript (finding the starting point on a page may take several minutes, even for skilled collators).

Since this task is focused on close inspection of a single manuscript rather than an information seeking or structuring task, I wanted to encourage at least a minimal engagement with the spatial aspects of the system. To achieve this, the proctor initialized the workspace in a configuration that was expected to be slightly inconvenient (the base text was displayed to the left of the manuscript page) and demonstrated how to re-arrange the panels during the introduction to the tool if they should so desire.

During the course of the study, the proctor observed the participants during both tasks, using the final interview portion of the study to follow-up on any questions that arose about specific actions that a participant may have performed. The computer-based tasks were recorded using Camtasia and screen captures were taken of the final state of the workspace. Marked copies of the base text were kept with the records for each participant. Following the study, these materials were analyzed to determine numeric performance metrics including how many lines of the manuscript were collated, the number of variants correctly identified and the ratio of correctly identified variants to the total number of variants (correct, missed, and incorrectly identified).

Following each task, participants indicated whether they disagreed or agreed on a scale of 1 to 5 (1: strongly disagree, 5: strongly agree) that they were confident in the accuracy of their collation, that the collation process was straightforward and that they encountered problems during the collation. The post-task survey also included a space for general comments about the task. This approach to data collection was intended blends direct observation, semi-structured interviews and numerical performance metrics to yield both quantitative and qualitative data.

5.3.2 Participants

The participants came with a wide range of backgrounds[7]. Most were students at the Seminary in either the M.Div program or the Ph.D. program. 12 participants were current or former staff members with CNTTS, 5 were faculty members at NOBTS or other universities in the area, 4 were not affiliated with the center, but had taken course work in textual criticism. Each participant was asked to rate his or her skill as a collator on a scale of 1 (Expert) to 4 (Novice).

The base demographics may be misleading. Faculty, for instance, may have had little training in collation or may have not performed collations recently. The visiting scholar rated himself a novice (4) in terms of experience since this was the first time he had conducted a collation of a manuscript. However, he had completed a Ph.D. project in textual criticism and had extensive experience working with manuscripts.

---

[7] Appendix A provides a complete listing of all participants.

Correspondingly, his collation was quite good when compared to more experienced participants.

5.3.3 Results and Observations: Supporting Collation

I present and discuss the results of this study in two parts. This section covers the results specific to supporting the manuscript collation task. The following section looks at the use of the spatial features of CritSpaceTC.

Two measures are important for assessing a tool to support collation: speed and accuracy. Collation is a time-intensive process and, consequently, the speed of collation (as measured by the number of line in the manuscript that a collator was able to examine) is a major factor in selecting whether to adopt a new tool or procedure. Accuracy, however, is far more important. No speed improvements can justify a significant decrease in accuracy.

The performance results are summarized in table 1. While participants were about as accurate using CritSpaceTC for their collations as they were on paper, they were significantly slower. In terms of accuracy, when using the computer-based collation, the ratio of correctly identified variants to all variants (both missed and incorrectly identified) was 0.9 (stdev: 0.10%) compared to 0.87 (stdev: 0.14) for the paper based collations; a statistical tie ($p=0.46$). Participants were able to collate significantly more manuscript lines on paper than they were on the computer. Excluding

Table 1. Summary of collation performance

	Computer Based Collation (Chpt. 20)					Computer Based Collation (Chpt. 22)				
	Lines	Variants Correct	Variants Missed	Variants Added	Ratio	Lines	Variants Correct	Variants Missed	Variants Added	Ratio
Avg	28.25	15.13	0.50	0.38	0.93	43.00	13.50	1.63	0.13	0.87
Stdev	11.56	5.64	0.76	0.52	0.09	17.23	8.26	1.60	0.35	0.11

	Paper Based Collation (Chpt. 20)					Paper Based Collation (Chpt. 22)				
	Lines	Variants Correct	Variants Missed	Variants Added	Ratio	Lines	Variants Correct	Variants Missed	Variants Added	Ratio
Avg	50.38	20.88	1.75	0.75	0.90	49.75	16.13	1.50	0.25	0.83
Stdev	33.27	10.89	2.19	0.71	0.09	20.47	10.73	1.04	0.46	0.18

data from two participants with anomalous paper-based collations[8], participants collated an average of 54.43 (stdev: 25.64) lines when using paper and 34.86 (stdev: 17.06) lines when using CritSpaceTC. These results are statistically significant (p=0.025).

The reported survey data yielded similar results and corresponds well to the performance metrics. Participants reported a high degree of confidence in their accuracy with no statistically significant difference between the paper and computer-based tasks. They reported that the process of collations less straightforward and more strongly agreed that they encountered problems using the computer as compared to paper-based collations (both statistically significant).

[8] Subj 02 started in the wrong location, Subj 13 had to leave the paper based study early

A number of usability issues were identified over the course of the study and the results need to be interpreted in light of those issues as they likely had a significant impact on both the speed of collation and user perceptions about their work and the enjoy-ability of using the system. The CritSpaceTC system was missing a few features. Most notably, the ability to represent transpositions (a re-ordering of the words in the text) explicitly rather than simply retyping the text as it appears. For NT scholars, transposed text and changed text are fundamentally different mental categories. Users (unsurprisingly) found that the keyboard commands for opening the window to mark variant text cryptic and would have preferred a button on the toolbar. In addition to these feature oriented problems, a number of bugs were discovered during the test including both relatively minor glitches (one, for example, caused a variant to be displayed incorrectly) and a recurring fatal error that occurred with about four users who lost their session when the page reset to the browser's home-page. These issues resulted in a system that felt rough around the edges and adversely impacted both collation speed and the users' experience. The paper-based interface had a bug as well, as several experienced collators commented on the poor quality of the paper-based printouts provided by CNTTS for the study. If they had a choice, they would have gone back and made higher-quality print outs. These difficulties aside, the study's results support three broad conclusions.

First, collation with the computer is about as accurate as paper-based collation, even given the prototype status of this study and the fact that participants were working with this tool for the first time. While prior experiences at CNTTS have led them to have

serious reservations about collation support software yielding lower accuracy, this has been based on tools that used software to record variation, but relied on paper copies of the source manuscripts. Making the manuscripts available digitally provides users with the ability to zoom in to see additional detail. Participants reported that this made it much easier to read the digital text. In addition, participants reported that the contrast was better for the digital texts than for the printed text. In practice, accuracy for both the digital and paper-based versions of the task was very high, despite a fairly difficult text to collate.

Second, collation with the computer is slower than paper-based collation. This finding is expected—the participants were very familiar with paper-based collation and it is difficult to see how the mechanics involved in marking the variants using a computer interface can be made easier than making a paper base text. Nevertheless, the magnitude of the difference is significant (a 36% decrease in the number of lines examined). A number of factors, however, mitigate the ultimate impact of this slowdown.

Speed of collation using CritSpaceTC was adversely impacted by a number of factors including: bugs interfered with the workflow, the keyboard commands were difficult to learn, and the approach was fundamentally new to people. Subject 16 voiced a common theme in the interview, stating that the "computer is harder because it is new" and suggesting that he needed more practice. Even with those difficulties, he felt that it was a easier to collate from the computer because "the text is clear and you have the ability to manipulate the image." Many of the system level drawbacks that impacted

collation speed result from the prototype status of the application and are readily fixable in a more polished tool. More experience with the tool should also help improve speed.

In addition, as several participants pointed out, focusing on the speed of the initial collation misses the more important question of how long it takes to process a manuscript. Since collation is only one step in the process that involves two independent collations, the reconciliation of discrepancy and ultimately updating the apparatus), gains in the later stages may possibly offset the impact on speed during the collation stage. Having the data recorded in the computer could facilitate comparison of multiple collations and make the process of entering this data into the final critical apparatus faster and more accurate.

Third and finally, the study delivered mixed results with respect to the users' experience. The survey data showed that participants clearly preferred paper-based collation to computer-based collation, but the follow up interviews show potential for this opinion to switch in favor of computer-based collation. Some aspects of the collation process were clearly more difficult to perform using the computer. Most notably, participants found it difficult to keep their place on the base manuscript during collation. In the print based collation, right-handed users would mark the base text with their right hand, while following along in the manuscript with their left hand. This allowed them to keep one finger on the last word that they had read while they focused on recording the variant on the base text. While a number of participants used their finger as a guide in reading the manuscript on the computer screen, marking a variant required that they place both hands on the keyboard. This frequently caused them to lose

their place and slowed their progress. A few users devised strategies or suggested enhancements to overcome this limitation that will be discussed in the following section.

On the other hand, many participants seemed to genuinely enjoy using CritSpaceTC. The words of subject 4, "I really surprisingly enjoyed using the computer" were echoed by several participants. This enjoyment increased as participants became more familiar with the interface, and a number of people indicated that they felt that they had just gotten the hang of it when the time ended. For many, their appreciation for CritSpaceTC was stated generally or else stemmed from the higher image quality. Several of the more experienced collators, however, offered an important insight as to why this may be: paper-based collation hurts. A collator may spend hours at a time hunched over a manuscript, straining his eyes to read ancient characters and to distinguish between historical damage, stray marks and ink. By moving this work to a computer-screen, subject 15 found that the process was "less physically and mentally straining." This is a particularly important finding since it indicates that, over the course of more long-term collation work, the ergonomic improvements of the computer based interface may ultimately result in improved speed and accuracy as collators are able to work longer with less strain.

In summary, CritSpaceTC was able to support a task that several at CNTTS initially thought could not be performed on a computer to the same level of quality as it was on paper. Despite the rough edges, participants were just as accurate using CritSpaceTC as they were using paper. The drawbacks in terms of speed may persist

even in the face of an improved implementation, but the speed gap should narrow. Overall, CritSpaceTC proved to be a promising platform for supporting textual criticism.

5.3.4 Results and Observations: Spatial Interactions

While evaluating the success of CritSpaceTC in supporting textual criticism (TC) is straightforward and direct (measuring speed, accuracy and user preferences), evaluating the utility of spatial interacts is more nuanced and indirect. In order to avoid encouraging participants to use the spatial features in ways that were artificial and unlikely to be useful in more realistic settings, the study intentionally downplayed these features. Participants were told that they could rearrange the workspace and resize the panel (to be sure that they were aware of these features) and the initial setup of the space encourages a basic rearrangement. Beyond that, there was nothing in the study setup or design that was done to encourage interaction with the spatial features.

My goal in evaluating the use of spatial features was to identify distinctive individual practices rather the than statistically generalizable commonalities used to demonstrate support for TC. These deviations from usual behavior are not outliers to be discarded. Instead they offer insights into how CritSpace facilitates idiosyncratic customizations that support personalized use and novel tweaks that help optimize the workflow or overcome technological difficulties. Where future use subsequently proves these deviations to be broadly applicable, they can be seen as innovations that will spread to become more widely used within the community.

For the most part, people did not pay much attention to the spatial features of the system. The post task interviews focus on the support (or lack thereof) for textual criticism. This indicates that, at a minimum, the spatial features did not get in the way. Being unremarkable (without being unused), in this case, is a good sign.

Most participants rearranged the workspace to place the image on the left hand side of the screen and the base text on the right hand side in a layout that mimics the arrangement used in the paper-based workflow. This is unsurprising in itself since this layout was suggested as part of the introduction to CritSpaceTC. More significant is that a few of the participants didn't wait until that part of the introduction; upon seeing the image initially placed to the right of the base text, they immediately suggested that it should be the other way around. While participants may not have commented on the support for re-arranging their workspace, the layout was clearly something that was important to them.

While most participants used roughly the same workspace arrangement, simply reversing the order of the panels, a few gave more thought to how this arrangement could be adjusted might support their work. These arrangements were often used to help mitigate the difficulty of keeping track of their location in the base text. One re-sized the image panel so that it showed only one column of text at a time. Another adjusted the zoom level so that the line height on the manuscript matched the line height on the base-text. Yet another shortened the image panel so that it only displayed three lines of text. One participant found that if he resized the base text display, the line wrapping would closely match that of the manuscript, making it easier to find his place in the text. All of

these different strategies represent relatively minor, *ad hoc* tweaks to the display that made life easier for individual collators.

One participant, a former CNTTS staff member who is currently working on a different doctoral project, discussed some of the opportunities that he saw for the system to support his current research. His work involves comparing multiple images from different manuscripts in order to study changes in orthography over time. He suggested pulling several different page images into the workspace so that they could be easily compared.

In summary, participants made use of CritSpaceTC's facilities in simple, straightforward ways and were far more interested in talking about with me about textual criticism than about the spatial features of the system. While most did not take advantage of the spatial features beyond re-arranging the panels, a few discovered that they could use these features to make their work easier. One commented on the possibilities for this CritSpaceTC to be tweaked to support a different manuscript analysis task that he was working on. These findings support the conclusion that the spatial interactions provided by CritSpace are useful tools for supporting the day-to-day use of source documents by cultural heritage scholars, even in tasks where the primary focus is not gathering and organizing information. They largely recede into the background but are present and discoverable when needed.

It is worth emphasizing that these results were observed within the context of a 5 minute introduction to the system followed by 45 minutes of use. As scholars have the ability to spend more time with the system and more leisure to experiment, it is

reasonable to expect that innovation will continue and successful innovations will be shared within the user community.

5.4 Summary

This evaluation provides a window onto a full-cycle customization, deployment and use of the CritSpace framework within a specific application domain. It demonstrated the usefulness of key design decisions such as the ability to integrate data from multiple sources, the reuse of existing JavaScript components in the development of interactive panels, and the usefulness of rich support for creating behavior interactions between panels. The TC domain supported by this task, unlike much of the previous work with interactive visual interfaces, emphasized intensive, simultaneous of content presented in different panels. The CritSpaceTC system proved to be an effective tool for supporting manuscript collation (albeit, one that is still in need of polishing before it is ready for deployment in production). And despite the fact that the task used for the summative evaluation focused on intensive analysis rather than knowledge organization, the study participants found innovative uses for the spatial interactions afforded by the system.

CHAPTER VI

CONCLUSIONS AND FUTURE WORK

6.1 Conclusions

This work makes two significant contributions to the body of knowledge. First, it responds to need for in-depth studies on the research and work practices within humanities scholarship and how to effectively design user-interfaces that support those practices. Both the formative (Chapter II) and summative (Chapter V) studies contribute to a deeper understanding of what cultural heritage scholars actually do when they do research. This is a desperately under-studied area.

Secondly, my work offers insight into how to design and build systems that provide both lightweight, flexible interfaces that avoid requiring users to commit to explicit formalization of knowledge, while at the same time provided carefully designed support for domain specific interactions. Digital libraries contain richly structured data and scholarly disciplines develop semi well-defined work practices for using that data. Traditional approaches to spatial hypertext systems do not make effective use of this existing structure because they were designed to support different tasks. CritSpace offers a both/and solution: lightweight *ad hoc* interactions to support open ended exploration of new knowledge along with custom support for well-defined tasks.

Finally, my work makes a practical contribution in the form of the CritSpace system itself. This application provides a general purpose platform for building rich, web-based interactive visual interfaces. These systems can be integrated with existing

digital libraries relatively easily. Consequently, I expect CritSpace to serve as a useful platform for both research and production environments.

6.2 Future Work

The work presented here focuses on work characterized by in-depth analysis. Scholars could also benefit from CritSpace as an interface for knowledge seeking and organization. One possible direction is to deploy CritSpace as a front-end to a widely used digital humanities repository such as NINES [2011]. NINES presents aggregated proxy representations of digital objects from more than 100 scholarly projects in an interface that allows users to create private collections and exhibit those collections. Deploying CritSpace in this context would enable evaluation with a large user community and a wider range of research needs and tasks. Since NINES is actively used by scholars around the US and the world, a multi-stage evaluation that involved automatically collected system-usage data, traditional usability studies, surveys and interviews.

Prior work in the field of spatial hypertext has emphasized computational support to recognize implicit structure in the way objects are laid out in the workspace [Shipman et al. 1995] and in the use of history mechanisms to provide users with a way to replay the knowledge organization process [Shipman et al. 2001a]. Addressing these topics in the context of CritSpace raises a number of interesting questions. For example, when working with externally defined objects (digital facsimiles, for example), how should a history mechanism interleave changes that take place within the CritSpace environment

and those that take place through an external tool such as a dedicated facsimile editor? Are there ways to integrate layout analysis algorithms that make use of specific characteristics of the target domain in order to make more intelligent inferences or to discover domain specific knowledge structures?

One motivation for the framework-based design that I have adopted with CritSpace is my desire to see this system serve as a general platform for research into how highly-interactive visual interfaces can be used to support humanities research in general and research that requires access to image-based data in particular. Making CritSpace more widely available and useful to the digital humanities community will require further development and documentation.

REFERENCES

AUDENAERT, N., AND FURUTA, R. 2008. Annotated facsimile editions: defining macro-level structure for image-based electronic editions. In *Proceedings of the Digital Humanities Conference (DH '08).* University of Oulu, Finland, 47-49.

AUDENAERT, N., LUCCHESE, G., SHERRICK, G., AND FURUTA, R. 2008. CritSpace: Using spatial hypertext to model visually complex documents. In *Proceedings of the Digital Humanities Conference (DH '08)*. University of Oulu, Finland, 50-53.

AUDENAERT, N., AND FURUTA, R. 2010. What humanists want: how scholars use source materials. In *Proceedings of the 10th Annual Joint Conference on Digital Libraries (JCDL '10).* ACM, New York, NY, 283-292.

AusLit: The Australian Literature Resource. http://www.austlit.edu.au/ (accessed August 2011)

BAE, S., KIM, D., MEINTANIS, K., MOORE, J.M., ZACCHI, A., SHIPMAN, F., HSIEH, H., AND MARSHALL, C. 2010. Supporting document triage via annotation-based multi-application visualizations. In *Proceedings of the 10th Annual Joint Conference on Digital libraries (JCDL '10)*. ACM, New York, NY, USA, 177-186.

BATES, M. J. 1996. The Getty end-user online searching project in the humanities: report no. 6: overview and conclusions. *College & Research Libraries, 57*, 6, 514-523.

BERNSTEIN, M. 2002. Storyspace 1. In *Proceedings of the 13th ACM Conference on Hypertext and Hypermedia*. ACM, New York, NY, 172-181.

BERNSTEIN, M. 2011. Can we talk about spatial hypertext? In *Proceedings of the 22nd ACM Conference on Hypertext and Hypermedia (HYPERTEXT '11).* ACM, New York, NY, 103-112.

BORGMAN, C.L. 2009. The future is now: a call to action for the digital humanities. *Digital Humanities Quarterly*, *3*, 4, 233.

BRADLEY, J., AND VETCH, P. 2007. Supporting annotation as a scholarly tool—experiences from the Online Chopin Variorum Edition. *Literary and Linguistic Computing 22*, 2, 225-241.

BRADLEY, J. 2008. Thinking about interpretation: Pliny and scholarship in the humanities. *Literary and Linguistic Computing 23*, 3, 263-279.

BROCKMAN, W.S., NEUMANN, L., PALMER, C.L., AND TIDLINE, T.J. 2001. Scholarly Work in the Humanities and the Evolving Information Environment, A report from the Council on Library and Information Resources.

BUCHANAN, G., BLANDFORD, A., THIMBLEBY, H., AND JONES, M. 2004. Integrating information seeking and structuring: exploring the role of spatial hypertext in a digital library. In *Proceedings of the 15th ACM Conference on Hypertext and Hypermedia (HYPERTEXT '04).* ACM, New York, NY, 225-234.

BUCHANAN, G., CUNNINGHAM, S.J., BLANDFORD, A., RIMMER, J., AND WARWICK, C. 2005. Information seeking by humanities scholars. In *European Conference on Digital Libraries (ECDL '05)*. Springer Verlag, Berlin, Germany, 218-229.

CARD, S.K., MACKINLAY, J.D., AND SHNEIDERMAN, B. (EDS.) 1999. *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufmann Publishers, San Francisco, CA, 1-34.

CHU., C. 1999. Literary Critics and their information needs: a research phase model. *Library and Information Science Research 21*, 2, 247-273.

CLEMENT, T., PLAISANT, C., AND VUILLEMOT, R. 2009. The Story of One: Humanity scholarship with visualization and text analysis. In *Proceedings of the Digital Humanities Conference (DH '09)*. Online at http://www.cs.umd.edu/localphp/hcil/tech-reports-search.php?number=2008-33 (accessed August 2011)

CODEX SINAITICUS PROJECT 2009. http://codexsinaiticus.org/en/ (accessed August 2011)

CONKLIN, J., AND BEGEMAN, M.L. 1988. gIBIS: a hypertext tool for exploratory policy discussion. *Transactions on Information Systems 6*, 4, 303-331.

CSIKSZENTMIHALYI, M. 1990. *Flow: The Psychology of Optimal Experiences*. Harper & Row, New York, NY.

DALTON, M.S. AND CHARNIGO, L. 2004. Historians and their information sources. *College & Research Libraries 31,* 2, 400-425.

DEROSE, S., DURAND, D., MYLONAS, E., AND RENEAR, A. 1990. What is Text Really? *Journal of Computing in Higher Education 1*, 2, 3-26.

DRUCKER, J. AND NOWVISKIE, B. 2004. Speculative Computing: Aesthetic Provocations in Humanities Computing. In S. Shcreibman, R. Siemens, and J. Unsworth (eds.), *The Blackwell Companion to Digital Humanities*, Blackwell Publishing Ltd, Oxford, UK, 431-447.

ELLIS, D. 1989. A behavioural model for information retrieval system design. *Journal of Information Science 15,* 4/5, 237-247.

KOCHUMMAN, R., MONROY, C., FURUTA, R., GOENKA, A., URBINA, E, AND MELGOZA, E. 2002. Towards an electronic variorum edition of Cervantes' Don Quixote: visualizations that support preparation. In *Proceedings of the 2nd Annual Joint Conference on Digital Libraries (JCDL '02)*. ACM, New York, NY, 199-200.

LAVAGNINO, J. 2006. When not to use TEI. In Burnard, L., O'Keeffe, K. O., Unsworth, J., (eds.) *Electronic Textual Editing*. New York: Modern Language Association. http://www.tei-c.org/Activities/ETE/Preview/lavagnino.xml (accessed 10 November 2006)

MARSHALL, C., HALASZ, F., ROGERS, R. AND JANSSEN, W. 1991. Aquanet: A hypertext tool to hold your knowledge in place. In *Proceedings of the Third ACM Conference on Hypertext (HYPERTEXT '91)*. ACM, New York, NY, 261-275.

MARSHALL, C., AND ROGERS, R. 1992. Two years before the mist: experiences with Aquanet. In *Proceedings of the ACM Conference on Hypertext (ECHT '92)*. ACM, New York, NY, 53-62.

MARSHALL, C. AND SHIPMAN, F. 1993. Searching for the missing link: Discovering implicit structure in spatial hypertext. In *Proceedings of the Fifth ACM Conference on Hypertext (HYPERTEXT '93)*. ACM, New York, NY, 217-230.

MARSHALL, C., SHIPMAN, F., AND COOMBS, J. 1994. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the ACM European Conference on Hypermedia Technology (ECHT '94)*. ACM, New York, NY, 13-23.

MARSHALL, C. AND SHIPMAN, F. 1997. Spatial hypertext and the practice of information triage. In *Proceedings of the Eighth ACM Conference on Hypertext (HYPERTEXT '97)*. ACM, New York, NY, 124-133.

MCGANN, J. 1996. The Rosetti Archive and image-based electronic editing. In Finneran, R. J. (ed.) *The Literary Text in the Digital Age*. U. of Michigan Press, Ann Arbor, MI, 145-183.

Microsoft OneNote 2010. http://office.microsoft.com/en-us/onenote/ (accessed August, 2011).

MindManager by Mindjet Mindjet web site: http://www.mindjet.com, (accessed August 2011).

NAKAKOJI, K., YAMAMOTO, Y., TAKADA, S., AND REEVES, B.N. 2000. Two-dimensional spatial position as a means for reflection in design. In *Proceedings of the 3rd Conference on Designing Interactive Systems (DIS '00)*, ACM, New York, NY, 145-154.

NINES: Networked Networked Infrastructure for Nineteenth-Century Electronic Scholarship. http://www.nines.org (accessed August 2011)

RIMMER, J., WARWICK, C., BLANDFORD, A., GOW, J. AND BUCHANAN, G. 2008. An examination of the physical and the digital qualities of humanities research. *Inf. Process. Manage. 44*, 3, 1374-1392.

ROTHFEDER, J. MANMATHA, R., RATH, T.M., 2006. Aligning transcripts to automatically segmented handwritten manuscripts. In *Proceedings of the 7th International Workshop on Document Analysis Systems (DAS 2006).* LNCS 3872. 84-95.

SCHOEN, D. A. 1983. *The Reflective Practitioner: How Professionals Think in Action.* Basic Books, New York, NY.

SHIPMAN, F., MARSHALL, C., AND MORAN, T. 1995. Finding and using implicit structure in human-organized spatial layouts of information. In *Proceedings of the SIGCHI*

Conference on Human Factors in Computing Systems (CHI '95). ACM, New York, NY, 346-353.

SHIPMAN, F. AND MARSHALL, C. 1999. Spatial hypertext: an alternative to navigational and semantic links. *ACM Computing Surveys 31,* 4.

SHIPMAN, F. AND MCCALL, R. 1999. Supporting incremental formalization with the hyper-object substrate. *ACM Transactions on Information Systems 17,* 2, ACM, New York, NY, 199-227.

SHIPMAN, F., AND MARSHALL, C. 1999. Formality considered harmful: experiences, emerging themes, and directions on the use of formal representations in interactive systems. *Computer Supported Cooperative Work 8*, 4, 333-352.

SHIPMAN, F., HSIEH, H., MALOOR, P. AND MOORE, M. 2001a. The Visual Knowledge Builder: A second generation spatial hypertext. In *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia (HYPERTEXT '01)*, ACM, New York, NY, 113-122.

SHIPMAN, F. AIRHART, R., HSIEH, H. MALOOR, P. MOORE J. M., AND SHAH D. 2001b. Visual and spatial communication and task organization using the Visual Knowledge Builder. In Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work (GROUP '01), ACM, New York, NY, 260-269.

SHIPMAN, F., HSIEH, H., MOORE, M. AND ZACCHI, A. 2004. Supporting personal collections across digital libraries in spatial hypertext. In *Proceedings of the ACM and IEEE Joint Conference on Digital Libraries (JCDL '04)*. ACM, New York, NY, 358-367.

SHNEIDERMAN, B. 2007. Creativity support tools: accelerating discovery and innovation. *Communications of the ACM 50*, 12, ACM, New York, NY, 20-32.

SPENCER, M., BORDALEJO, B., WANG, L., BARBROOK, A. C., MOONEY, L. R., ROBINSON, P., WARNOW, T., HOWE, C.J. 2003. Analyzing the order of items in manuscripts of The Canterbury Tales. *Computers and the Humanities 37*, 1, 97-109.

STONE, S. 1982. Humanities scholars: information needs and uses. *Journal of Documentation 38*, 4, 673-691.

TIBBO, H. R. 2002. Primarily history: historians and the search for primary source materials. In *Proceedings of the ACM and IEEE Joint Conference on Digital Libraries (JCDL '02)*. ACM, New York, NY, 1-10.

TRIGG, R. 1988. Guided tours and tabletops: tools for communicating in a hypertext environment. *ACM Transactions on Information Systems 6*, 4, 398-414.

TOMS, E. G. AND O'BRIEN, H. L. 2008. Understanding the information and communication technology needs of the e-humanist. *Journal of Documentation 64*, 1, 102-130.

UNSWORTH, J. 2000. Scholarly Primitives: what methods do humanities researchers have in common, and how might our tools reflect this? Paper given at symposium "Humanities Computing: formal methods, experimental practice", London, May 13, 2000, King's College, London. http://www.iath.virginia.edu/~jmu2m/Kings.5-00/primitives.html (viewed December 2006).

WARWICK, C., TERRAS, M., GALINA, I., HUNTINGTON, P., AND PAPPA, N. 2008. Library and information resources and users of digital resources in the humanities. *Library and Information Resources 42*, 1, 5-27.

WATSON-BOONE, R. 1994. The information needs and habits of humanities scholars. *Reference Quarterly 34*, 2, 203-216.

Women Writers Project. Julia Flanders, Dir. http://www.wwp.brown.edu/ (accessed August 2011)

YUI. 2001. YAHOO! User Interface, Rich Text Editor. http://developer.yahoo.com/yui/editor/ (accessed August 2011)

YAMAMOTO, Y., NAKAKOJI, K., AND AOKI, A. 2002. Spatial hypertext for linear-information authoring: interaction design and system development based on the ARD design principle. In *Proceedings of the 13th ACM Conference on Hypertext and Hypermedia (HYPERTEXT '02)*, ACM, New York, NY, 35-44.

APPENDIX A

STUDY PARTICIPANTS

Group 1A	Sex	Age	Role	Coursework	Training	Time	Skill
Subj 04	M	22-36	Student	TC of GNT[9]	Work in CNTTS	0 yrs	3
Subj 05	M	27-35	Project Director RA, PhD Student	TC of GNT (MA & PhD)	Work in CNTTS	3 yrs	1
Subj 12	M	22-26	M.Div Student	TC of GNT	None	0 yrs	4
Subj 16	F	35-45	Graduate, Assoc. Prof of NT & Hist Theo	TC of GNT (PhD)		2 yrs	2
Subj 20	M	35-45	RA, PhD Student	TC of GNT (MA)	Work in CNTTS	2yrs	2

Group 1B	Sex	Age	Role	Coursework	Training	Time	Skill
Subj 08	M		Visiting Scholar	Ph.D in TC Project	First collation of an MS	0 yrs	4
Subj 10	M	45+	Professor	TC, Collation activity (years ago)	none	0 yrs	3
Subj 11	M	35-45	RA	TC of GNT, (MA & PhD)	Work in CNTTS some teaching	8 yrs	1
Subj 19	M	45+	Student	TC of GNT	none	0 yrs	4
Subj 21	M	45+	PhD. Candidate, RA from 2006-2010	MS Studies	Work in CNTTS	4 yrs	1

Group 2A	Sex	Age	Role	Coursework	Training	Time	Skill
Subj 02	M	22-26	M. Div, former Secretary	TC of GNT	Work in CNTTS	0.5 yrs	2
Subj 03	M	27-35	Ph.D., TA, RA	TC of GNT (MA, Ph.D)	Work in CNTTS	2.5 yrs	1
Subj 07	M	45+	M. Div	Undergrad. coursework	none	0 yrs	3
Subj 15	M	22-26	Secretary, RA	3 courses in Greek	Work in CNTTS	0.67 yrs	1
Subj 17	M	35-45	Ph.D. Student, Adjunct Prof.,	TC of GNT (MA, Ph.D)	Work in CNTTS	6 yrs	1

[9] Textual Criticism of the Greek New Testament. This class is offered at an introductory level at the Masters level (for students who have already completed some Greek language course work and as a seminar class at the P.hD. level.

Group 2B	Sex	Age	Role	Coursework	Training	Time	Skill
Subj 01	M	27-35	RA/Student	TC of GNT	Work in CNTTS	2 yrs	2
Subj 09	M	22-26	RA/Ph.D.	TC of GNT (MA, Ph.D)	Work in CNTTS	2 yrs	2
Subj 13	M	35-45	Assoc Prof of GNT, Regional Dean for the Extension Center System	Ph.D. in NT, Doctorate in TC	none	??	1
Subj 14	F	22-26	MA	TC of GNT	none	0 yrs	4
Subj 18	M	45+	Prof. of NT	none	IGNTP Training	0.5-1 yr	2

VITA

Michael Neal Audenaert received his Bachelor of Science degree in computer science in 2001, and his Master of Science in 2006 from Texas A&M University. During his graduate studies he has served on the Graduate Student Council and as President of the Computer Science Graduate Student Association. He served as a member of the Texas A&M University Digital Humanities Working Group and is the continuing issues editor for the Bulletin of the IEEE Technical Committee on Digital Libraries. His research interests focus on digital library interface, digital humanities, and pattern classification. He has published 14 peer-reviewed articles and papers in these areas.

In 2009, Neal founded the Institute for Digital Christian Heritage (IDCH), a non-profit organization dedicated preserving Christianity's history and making it accessible to new audiences.

Neal can be reached at Institute for Digital Christian Heritage, 1700 George Bush Dr. E., Ste 220, College Station, TX, 77840. His email address is neal@idch.org.