

OPEN ANNOTATION COLLABORATION
PHASE 1: DATA MODEL & INTEROPERABILITY SPECIFICATION;
AXE / ZOTERO INTEGRATION; SCHOLARLY ANNOTATION ANALYSIS

Executive Summary

The University of Illinois at Urbana-Champaign, in partnership with the Los Alamos National Laboratory, George Mason University, the University of Queensland, and the University of Maryland, and in collaboration with JSTOR, requests support to be expended over fourteen months (16 May 2009 through 15 July 2010) in order to: 1) define an advanced, standards and practice-based, scholarly-focused framework for sharing, interoperating across and exploiting annotations of digital resources; and 2) lay the foundation for deploying this framework in subsequent phases of work. We will begin this effort with a systematic multi-perspective analysis of current annotation models, application designs and system architectures, performed in concert with an examination of a broad range of scholarly practices and scholarly-focused use cases involving annotations. This analysis will inform the development of a shared annotation data model supportive of interoperable annotations, adaptable by existing systems, and rooted in scholarly practice. In parallel, and also informing the definition and development of our shared, interoperable data model of scholarly annotation, we will integrate the Ajax XML Encoder (AXE) annotation libraries created by the Maryland Institute for Technology in the Humanities into the Zotero digital collection and citation management application created by the Center for History and New Media at George Mason University. This integration will provide Zotero with a new embedded scholarly annotation tool. The process of preparing for and doing this integration, in addition to informing the development of our interoperable annotation data model, ensures the ability to add interoperable annotation features and services to Zotero in a subsequent phase of work. The capstone deliverable of this initial project phase will be the public release with request for comments of an alpha-stage annotation interoperability specification, embedding our interoperable annotation data model and defining the read annotation interfaces required to implement this data model in practice.

Pending the success of this initial phase of work and the availability of additional funding, this grant will lay the foundation for subsequent implementation and deployment efforts, leading to the emergence of a ubiquitous Web and Resource-centric interoperable annotation environment that allows leveraging annotations across the boundaries of annotation clients, annotation servers, and content collections. This work cannot be done in a vacuum and must be fully cognizant of scholarly context, existing applications and real-world technical environments. Any new standard must make sense in the context of scholarly needs, dominant Web technologies, existing tools, and existing resource collections. The team assembled for this project includes scholars conversant with humanities scholarly practices -- both traditional print-based and emergent digital-based, technologists experienced in the successful creation of standards and data models conformant to the Web Architecture, repository and collection managers familiar with a range of media types and formats, and application developers with a proven track record for creating useful applications, tools, and services. Ultimately, our overriding goal is to advance the quality and functionality of widely available annotation applications as a way to encourage and facilitate use of digital resources by humanities scholars, while simultaneously enhancing the immediate and potential long-term value of annotations as Web resources in their own right by providing a framework and a set of rules for sharing them across annotation applications.

OPEN ANNOTATION COLLABORATION (OAC):
ENABLING INTEROPERABLE ANNOTATION OF SCHOLARLY DIGITAL RESOURCES

**PHASE 1: DATA MODEL & INTEROPERABILITY SPECIFICATION;
AXE / ZOTERO INTEGRATION; SCHOLARLY ANNOTATION ANALYSIS**

Co-Principle Investigators:

Timothy W. Cole	<i>University Library and Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign</i>
Neil Fraistat,	<i>Maryland Institute for Technology in the Humanities, University of Maryland</i>
Jane Hunter	<i>eResearch Laboratory, School of Information Technology & Electrical Engineering, University of Queensland</i>
Herbert Van de Sompel	<i>Research Library, Los Alamos National Laboratory</i>

Additional Co-Investigators:

Daniel Cohen	<i>Center for History and New Media, George Mason University</i>
Carole Palmer	<i>Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign</i>
Allen Renear	<i>Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign</i>
Robert Sanderson	<i>Research Library, Los Alamos National Laboratory (planned)</i>
John Burns ‡	<i>JSTOR</i>
Clare Llewellyn ‡	<i>JSTOR</i>

‡ Gratis ad hoc participation; affiliated with an institution not receiving budget as part of this project.

Table of Contents

Executive Summary	1
I. Goals and Objectives	4
<i>Project Objectives: Phase 1</i>	4
<i>Project Objectives: Subsequent Phases (separate funding)</i>	6
II. Impact.....	7
<i>Significance</i>	7
<i>Benefits to Content Providers</i>	8
<i>Benefits to Scholars</i>	9
III. State of the Art	11
IV. Statement of Work	14
Thread 1: Annotation Interoperability Specification	16
Thread 2: Implementation and Integration Experience.....	20
Thread 3: Scholarly Annotation Application Analysis & Model Evaluation	21

I. Goals and Objectives

The overarching goals of this project (consisting of multiple phases) are:

- To facilitate the emergence of a Web and Resource-centric interoperable annotation environment that allows leveraging annotations across the boundaries of annotation clients, annotation servers, and content collections. To this end, interoperability specifications will be devised.
- To demonstrate through implementations an interoperable annotation environment enabled by the interoperability specifications in settings characterized by a variety of annotation client/server environments, content collections, and scholarly use cases.
- To seed widespread adoption by deploying robust, production-quality applications conformant with the interoperable annotation environment in ubiquitous and specialized services and tools used by scholars (e.g., JSTOR, Zotero, and MONK).

A shared Annotation Data Model and a set of Read Annotation Interfaces will enable sharing annotations beyond the boundaries of individual solutions or content collections, and hence will allow for the emergence of value-added cross-environment annotation services. They also will facilitate the implementation of advanced end-user annotation services targeted at humanities scholars and capable of operating across a broad range of both scholarly and general collections. Further, they will enable customization of annotation services for specific scholarly communities, without reducing interoperability. The proposed work also will enable more robust machine-to-machine interactions and automated analysis, aggregation and reasoning over distributed annotations and annotated resources. By grounding our work in a thorough understanding of Web-centric interoperability and embedded models implemented by existing digital annotation tools and services, we will create an interoperable annotation environment that will allow scholars and tool-builders to leverage prior tool development work and traditional models of scholarly annotation, while simultaneously enabling the evolution of these models and tools to make the most of the potential offered by the Web environment.

Project Objectives: Phase 1

The four primary outcomes of phase 1 of this project will be:

- An analysis of existing annotation models, systems, and architectures and scholarly practices regarding annotations in order to inform the development of a shared annotation data model supportive of interoperable annotations, adaptable by existing systems, and rooted in scholarly practice.
- A publicly released first (alpha) annotation interoperability specification consisting of an Annotation Data Model and Read Annotation Interfaces.
- The successful integration of Ajax XML Encoder (AXE) code libraries (Maryland Institute for Technology in the Humanities [MITH]) into Zotero (Center for History and New Media [CHNM]), This integration will provide Zotero with a new, embedded scholarly annotation tool. In addition to advancing the development of both AXE and

Zotero, the process of preparing for and doing this integration will inform the development of our interoperable annotation data model and ensure our ability to add interoperable annotation features to Zotero in a future phase of this project.

- A complete or nearly complete first-draft minimal reference implementation, demonstrating proof-of-concept feasibility and laying a foundation for a follow-on phase of work (separately funded) featuring demonstrations and reference implementations exploiting real-world repositories such as *JSTOR*, *Flickr Commons*, and *MONK* (University of Illinois at UC), and leveraging existing scholarly annotation applications such as the new Zotero tool described above, *Pliny* (King's College) and *Co-Annotea* (University of Queensland).

Realistic, Concrete Use Case Scenarios

To be generally useful, development of the Annotation Data Model and Read Annotation Interfaces must be grounded both in existing knowledge regarding annotations and existing Web-centric interoperability approaches. In order to be useful for scholarly annotation, the data model and interoperability approach must be cognizant of actual scholarly practice and must accommodate a broad range of media types and repository sophistication.

As a means to help ensure concreteness and applicability of data model and interoperability specification developed, we have preliminarily identified seven potential annotation use case scenarios that once further elaborated will help guide our work. These scenarios will be refined, augmented and periodically revisited over the course of phase 1 of this project. An outcome of the scholarly annotation research and analysis undertaken as part of phase 1 of this project will be an enriched set of concrete, potentially demonstrable scholarly annotation use cases.

Contingent on success of the work in phase 1 and the availability of additional funding, we would propose in subsequent phase(s) of this project to demonstrate and create reference implementations around these or similar use cases:

- Consistent with the nascent annotation tool that will emerge as we integrate MITH AXE annotation libraries into Zotero, we will define an annotation scenario involving citation linking into performance videos from at least three different video archives of varying technical sophistication -- each with differing facility for addressing discrete segments within video content served. This will illustrate the flexibility to support differing precision of annotation referencing as appropriate to repository capability.
- Consistent with the nascent annotation tool that will emerge as we integrate MITH AXE annotation libraries into Zotero, we will define an annotation use care scenario that exercises the ability to use streaming audio to annotate streaming video, and to do so in a manner that allows subsequent merger of audio commentary stream with annotated video stream by an independent application.
- Consistent with the nascent annotation tool that will emerge as we integrate MITH AXE annotation libraries into Zotero, we will define an annotation use care scenario that exercises the ability to annotate text in a manner that protects confidentiality of annotation sources and targets, but still enables human-initiated and monitored machine-to-machine annotation sharing.
- Extrapolating from data models implicit to the current Co-Annotea and nascent Zotero

annotation tools, we will define an annotation use case scenario featuring a brokering service to extract social tagging annotations of digital still images held in Flickr Commons and del.icio.us. Successful demonstration (in this later phase of work) of the ability to harvest, aggregate, sort / merge / rank and present heterogeneous annotations of same or related targets held in multiple repositories, and to access them using Co-Annotea and Zotero will illustrate a key benefit of annotation interoperability.

- Building on the LORE (Literary Object Re-use and Exchange) annotation authoring reference implementation we will define an annotation use case scenario that enables users to annotate multiple mixed-media objects and the relationships between them in order to create compound objects. Such functionality would then allow objects to be published in accord with the shared data model and interoperability specification so as to enable annotations to be subsequently discovered, re-used, and enhanced (i.e., demonstrating ability to annotate annotations and annotation relationships).
- Extrapolating from an anticipated Pliny-based interoperable annotation reference implementation, we will define an annotation use case scenario that enables users to assert new annotation relationships between existing digital resources, such as between literary analysis in JSTOR digitized journal articles, digitized novels in MONK, and digitized history texts held by the Open Content Alliance. Demonstrating such functionality would illustrate the emerging scholarly practice of *netchaining*. As part of this scenario we propose demonstrating how annotation sharing enables *netchains* to be extended collaboratively.
- Extrapolating from an anticipated Pliny-based interoperable annotation reference implementation embedding the Pliny concept of *annotation containers*, we will define an annotation use case scenario to demonstrate the benefits of an annotation data model that supports compound annotation targets and sources -- e.g., targets or sources comprised of arbitrary groupings or aggregations of Web resources. By being able to simultaneously link and annotate multiple discrete resources, users could exploit such a data model to support scholarly practice primitives such as *comparing* and *illustrating*.

Bring cognizant from the start of concrete scholarly use cases and working in close collaboration with potential adopters, both content and service providers, will provide essential reality checks and ensure an initial specification draft that can be a viable starting point for subsequent implementation and that can evolve with application development and deployment. This in turn will allow us in future work to demonstrate the potential of an interoperable annotation environment, and as such increase the chances of adoption of the interoperability specifications beyond our immediate collaboration. The choice of content providers outlined in the above scenarios also will allow us to work towards interoperability of annotations across a variety of media types including text, images and video.

Project Objectives: Subsequent Phases (separate funding)

To verify the feasibility of the initial data model and alpha-release annotation interoperability specification developed in phase 1, and contingent on success of phase 1 and availability of additional funding, we anticipate a follow-on proposal to demonstrate the benefits of the interoperable annotation environment enabled by the specification and deploy applications exploiting that specification. We anticipate accomplishing these objectives by adapting existing annotation client and server tools and deploying three reference implementations, one of which

will be based on the popular Zotero browser plug-in enhanced with the MITH AXE annotation client libraries. These reference implementations will then be used to annotate both ubiquitous (e.g., *JSTOR*, *Flickr Commons*) and specialized (e.g., *MONK*, *AustLit*, *IMLS DCC / DLF Aquifer*) content collections. The inclusion of more specialized tools and content will illustrate the potential benefits of an interoperable annotation environment for specialized scholarly use cases. Finally, if subsequent phase or phases of work are funded, we will demonstrate value-added services created on the basis of annotations aggregated from heterogeneous annotation client / server reference implementations.

The outcomes of subsequent phases of this project will be:

- A refined, robust, stable, production-quality annotation interoperability specification evolved from that developed in phase 1 of the project.
- Three or more production-quality annotation client/server reference implementations of the interoperability specification to inform and guide other developers not directly involved in this project -- thereby encouraging broader exploitation of the interoperability specification in other annotation applications and across additional scholarly content.
- Multiple demonstrations of the interoperability specification across a range of scholarly content formats and annotation applications to confirm (and advertise) its broad utility and benefits.

Ultimately, our overriding goal is to advance the quality and functionality of widely available annotation applications as a way to encourage and facilitate use of digital resources by humanities scholars, while simultaneously enhancing the immediate and potential long-term value of annotations as Web resources by providing a framework and a set of rules for sharing them across annotation applications.

II. Impact

Significance

A number of issues and perceived obstacles have slowed adoption and exploitation of digital resources by scholars. Many of these have to do with the ease with which faculty and students can transfer traditional practices of scholarship and pedagogy into a digital context and/or with the perception of how adequately, or inadequately, emerging digital tools and content can be exploited for scholarly ends.

Annotating is a practice core to scholarship. In his 2000 presentation at King's College London, John Unsworth identified annotating as a "scholarly primitive" and noted with concern the lack of progress towards applications that effectively supported the sharing of annotations on the Web. He identified six other scholarly primitives, several of which in certain instances can be facilitated through annotation (e.g., *referring*, *comparing*, and *illustrating*).¹ Annotating is a

¹ Unsworth, building on a theme from Aristotle, uses the term *scholarly primitives* "to refer to some basic functions common to scholarly activity across disciplines, over time, and independent of theoretical orientation." Unsworth, John. 2000. "Scholarly Primitives: what methods do humanities researchers have in common, and how

pervasive element of scholarly practice for both the humanist and the scientist. It is a method used by scholars to organize knowledge and facilitate the creation and sharing of new knowledge. It is used by individual scholars when reading as an aid to memory, to add commentary, and to classify. It can facilitate shared editing, scholarly collaboration, and pedagogy. Yet scholars remain dissatisfied with the options available for annotating many digital resources. The importance of annotating as a scholarly practice coupled with the real-world limitations of existing practices and tools supporting annotation of digital content has had a significant retarding effect on the growth of digital scholarship and the level of digital resource use by scholars.

Many of the limits of current practice have to do with the limited reach of even well-founded efforts in this domain. While some good progress has been made on individual tools and in the context of isolated content repositories, the lack of a shared data model of annotation and consensus on an interoperability standard supporting easy exchange of annotation records has undercut the impact of such efforts. Scholars tend not to adopt new tools readily if the information gathered into these tools cannot be subsequently migrated to other applications they use every day or shared with colleagues easily. They also hesitate to adopt when the use of a tool is limited to a single content repository. As a result, scholars continue to express disappointment over their inability to annotate digital resources satisfactorily and cite this shortcoming of working with digital content as one reason why they are disinclined to make greater use of digital content. A 2008 Arizona State University Library study of faculty experiences with e-books reports faculty perceptions that "print books are more versatile [than e-books] for taking notes and making annotations" as well as a perception that as compared to using a physical copy of a book it is not as easy to work with the presentation of ideas and concepts that span discontinuous parts of a digitized book.² Intuitively it should be the other way round. Computer technologies should be making it easier to perform these activities.

The status-quo is such that scholars wanting to annotate are often confronted with having to learn different annotation clients for different content repositories, have no easy way to integrate annotations made on different systems or created by colleagues using other tools, and are often limited to simplistic and constrained models of annotations. Frameworks for annotation reference are inconsistent, not coordinated, and frequently idiosyncratic, and the constituent elements of annotations typically are not exposed to the Web Architecture as discrete resources, making annotations as information objects of interest in their own right more difficult to study. (The history of scholarship demonstrates the need to support such study.) Scholars are frustrated with digital resources that cannot be exploited as they routinely exploit print resources, and providers of digital resources are frustrated that scholars prefer to continue using print rather than digital resources for want of better ways to annotate. A lack of robust, interoperable scholarly-oriented tools for annotating across heterogeneous repositories of digital content and difficulties sharing or migrating annotation records in a standard are factors diminishing interest in integrating digital resources into scholarship.

Benefits to Content Providers

The proposed work addresses head-on many of these foundational issues and problems with

might our tools reflect this?" Presented at *Humanities Computing: formal methods, experimental practice*, King's College, London, May 13, 2000. Available: <http://www3.isrl.uiuc.edu/~unsworth/Kings.5-00/primitives.html> (viewed 20 November 2008).

² Carlock, Danielle M. and Anali Maughan Perry. 2008. Exploring faculty experiences with e-books: a focus group, *Library Hi Tech* 26 (2): 244-254.

the current regime. By doing so, our work will beneficially impact both scholars and scholarly content providers and will pave the way for advances in annotation application implementation. Specific illustrations of this potential will be demonstrated as part of our proposed work in order to show specific, concrete benefits and encourage additional implementations. Summarized here are the broader, high-level benefits we anticipate.

Annotation applications tightly bound to individual content repositories tend to model and express annotations and annotation relationships idiosyncratically. Such repository-centric annotation applications also tend to conflate storage of annotations with storage of content being annotated. Annotation models can be constrained by repository content models designed with non-annotation content in mind. Annotations may or may not be addressable independent of the annotated content, and thus annotation sharing and interoperability is made more difficult.

As described in the discussion of Thread 1 (below), we will exploit the efficacy of a Web- and Resource-centric model of annotation and annotation interoperability. Annotations instantiated according to such a model leverage established standards and are easily shared. Promulgation of a Web- and Resource-centric annotation model will encourage annotation application developers to build more broadly deployable applications and content providers to rely on standards-based ways to enable and facilitate annotation. This approach reduces pressures on content providers to invent or maintain their own idiosyncratic annotation applications. Content providers that conform to the Web Architecture model immediately benefit without having to create their own, repository-specific annotation application. By including in our work options for the caching of Annotation Sources and Annotation Targets (see description of Thread 1), and for identifying fragments of resources as Annotation Sources and Annotation Targets, we still provide ways to usefully create and share annotations even over content held in repositories that are not able to commit to long-term persistence or that fail to provide URIs at ideal granularity for the content they contain.

Content providers providing access to retrospectively digitized secondary scholarly literature (e.g., *JSTOR*, *Project Muse*, *AustLit*, *SETIS (Sydney Electronic Text and Image Service)*) benefit especially from the decoupling of repository and annotation application. Because journal articles and other forms of secondary literature reference and comment on primary sources extensively, passages in journal articles often serve as the source (i.e., the content) of an annotation. However, historical journal literature is also a primary source for scholarly research, not to mention that the reading of journal literature generally can suggest ideas and comments best captured as annotations. As such, passages in journal articles are frequently also the targets of annotation. Modeling this range of annotation possibilities in a repository-centric fashion is demanding. In a Web and Resource centric approach to annotation the modeling is more straightforward. In the Web graph, there is no intrinsic difference between an annotation source and an annotation target, both are Web resources -- the distinction is in the assertion of relationships expressed in the annotation record itself. The annotation source and annotation target are independently addressable, and the annotation record itself also is stored independently. This simplifies implementation and decouples annotation implementation from the repository.

Benefits to Scholars

Many current annotation applications rely on overly limited, simplistic models of annotation suffer from a tendency to conflate annotation source and annotation target. As Martin Mueller

suggests, the common, intuitive model of the annotation source as always a simple, brief, unformatted paragraph-length text is "clearly inadequate for most scholarly purposes, and it is not good enough for pedagogical purposes either." Mueller goes on to highlight the paradox that this condition should be prevalent in computer-assisted scholarship. "In a digital environment the content of a note can be anything that can be represented as a bit stream: a note can be a written text, a sound clip, an image, and so forth. This is a great advance over the print medium, but from another perspective, annotation software is in practice oddly restrictive when it comes to the structure of written notes."³ By implementing an annotation data model that explicitly distinguishes between annotation source, annotation target, and the assertions of annotation as expressed in an annotation record, assuming all to be Web resources, we enable and encourage more advanced scholarly annotation systems that allow scholars to generate and share the kinds of complex annotations required to support advanced scholarly practices. This benefits the scholar who wants to use and integrate digital resources into his or her scholarship, and makes it more likely he or she will do just that.

The recognition that scholarly annotating may be done for a range of reasons also is critical to ensure wide adoption of any generalized approach to annotation. While Muller (op. cit.) suggests that we can consider annotations simply as information produced for readers or viewers of resources, this can encompass a wide range of use cases. Consider, for instance, an annotation made for pedagogical purposes -- a professor elaborating a point in the text for the benefit of his or her students who will later read the text. Or consider annotations created as a means to facilitate resource discovery by would-be readers (e.g., social tagging as annotation). John Bradley suggests other instances where simple note-taking annotations are only the first step towards scholarly interpretations, which in turn may be represented by more complex annotations (or annotation containers as implemented in Pliny) having complex, multi-component annotation sources and/or complex annotation targets.⁴ Annotations may be used as a way to instantiate relationships between and among resources in a particular interpretive framework -- a way of classifying or clustering resources for the convenience of the annotator or other users, or to assert a condition of similarity or difference. Automated tagging of a digitized novel for part of speech is a form of annotation designed to facilitate subsequent machine-mediated reading and analysis. In all these instances, annotations provide, as Mueller puts it, a form of "reliable, just-in-time knowledge" at point of discovery and/or use. A Web and resource centric annotation model supports needed granularity and functionality and is built on a ubiquitous, well-established, well researched foundation. This approach gives confidence to scholars that they can migrate and/or share annotations they create with other applications and with colleagues.

Finally, while the focus of the demonstrations and implementations undertaken as part of this project will be on digital humanities scholarship, we anticipate similar benefits to the broader scholarly community. Returning to Unsworth's 2000 King's College presentation, he noted that a search in Google in 2000 for two of his humanities scholarly primitives together, annotation and comparison, yielded numerous hits to do with life sciences scholarship, specifically to do with the Human Genome Project. Repeating the search in November 2008, this time in Google Scholar, returns an even greater preponderance of scholarly annotation illustrations from the scientific community, alongside links to humanities digital text analysis research, research focused on the automated annotation of still images, and research into the semi-automated

3 Mueller, Martin. 2006. "Goals and Design Principles for Digital Annotation in a Scholarly and Pedagogical Environment." Unpublished draft.

4 Bradley, John. 2008. "Pliny: A model for digital support of scholarship". Journal of Digital Information, 9 (26). Online at: <http://journals.tdl.org/jodi/article/view/209/198> (viewed 21 November 2008).

annotation of a corpus of German newspapers. Advances in annotation modeling and the definition of a specification for sharing scholarly annotations will have significant benefits to digital scholarship extending beyond solely the digital humanities domain.

III. State of the Art

There is a significant amount of existing work that needs to be discussed and considered when proposing any new research on the subject of annotations, and how the state of the art will be advanced beyond the already mighty corpus of scholarly analysis that has been presented. Discussed below are some of the highlights of annotation research over the past twenty years intended to demonstrate the depth and breadth of existing research and understanding.

From the moment the concept of hypermedia or hypertext emerged, people have been discussing the concepts of annotating resources. The connection is natural and obvious: if one can join two bits of text together with some relationship, then joining a short piece of text describing, commenting on, or otherwise relating it to a portion of an existing text readily comes to mind. Both personal and collaborative note taking are among the early explorations of the use of annotations in hypermedia, from systems like QUILT⁵ and MINOS⁶, to Intermedia⁷, Xerox's NoteCards⁸ and Apple's HyperCard.

An early paper explicitly covering hypermedia annotation, as opposed to systems that support annotation-like behaviors, was published 15 years ago at the ACM conference on Hypertext, entitled, creatively, "The Knowledge Weasel Hypermedia Annotation System"⁹, in which the authors discuss the capabilities of their system that was built on the basis of free and publicly available software: a common format for representing annotations across media types, and support for collecting sets of annotations into groups. Much of the proposed model and the associated research agenda remains valid today, even though applications and software have moved inexorably onwards.

Catherine Marshall, then at Texas A&M University (and now at Microsoft Research), is a pioneer regarding annotations in digital libraries, publishing several highly cited papers in the mid-nineties, including "Annotation: From Paper Books to the Digital Library"¹⁰ which lays out an agenda of how academics (in this case students) use annotations on physical media and how their practice should be transferred into the digital realm. Generally, for research regarding how scholars use annotations, and how annotations can be best used in hypermedia, Marshall's published research will serve the project well.

5 Leland, M., Fish, R., Kraut, R. "Collaborative document production using quilt", *Procs of the 1988 ACM Conference on Computer-supported cooperative work*, pp 206-215, 1988

6 Christodoulaikis et al., "Multimedia document presentation, information extraction and document formation in MINOS: a model and a system", *ACM Transactions on Information Systems*, Vol 4, Issue 4, pp 345-383, 1986

7 Catlin, T., Bush, P., Yankelovich N., "InterNote: Extending a hypermedia framework to support annotative collaboration", *Procs of the second annual ACM conference on Hypertext*, pp 365-378, 1989

8 Monty, M., Moran, T., "A Longitudinal Study of Authoring Using NoteCards", *SIGCHI Bulletin*, Vol 18, Issue 2, 1986

9 Lawton, D., Smith, I., "The Knowledge Weasel Hypermedia Annotation System", *Procs of the Fifth ACM Conference on Hypertext*, pp 106-117, 1993

10 Marshall, Catherine C. "Annotation: from paper books to the digital library" *DL '97: Proceedings of the second ACM international conference on Digital libraries*, pp 131-140, 1997

Another extensively published author on annotation is Maristella Agosti at the University of Padova with 17 publications in the last five years on annotations in the realm of digital libraries and medieval manuscripts. In particular, a recent article “A Formal Model of Annotations of Digital Content”¹¹ is of relevance to the current work. In it, Agosti and Ferro describe a formalized model for annotation, which differs significantly from the initial model discussed above. In their model, for example, an annotation can only annotate a single object, but may be related to multiple documents in other ways. Our view is that this, while it may simplify the mathematical construction of formal proofs, is not a necessary constraint and the ability to have one annotation annotate multiple objects is a very valuable capability, for example supporting annotating depictions of the same person in multiple photographs.

Agosti’s is one of many models, each of which will need to be analyzed and assimilated into the overall model produced for interoperability. Possibly the best known Web-oriented annotation model is the RDF-based Annotea¹², produced by Kahan and Koivunen. Annotea has existed since 2001, but has never gained wide adoption despite its relatively sound model. This model will provide inspiration for the proposed project. Similar to Agosti’s model, the Annotea model allows only one target per annotation. In addition, Annotea does not treat an annotation as a resource in its own right, it does not distinguish between an annotation and the source resource of an annotation (see above), and it requires the content of the annotation (source resource in the above) to be created at the same time as the “annotates” relationship to the single target resource. These modeling choices prevent an annotation author from asserting the “annotates” relationship between two (or more) pre-existing resources.

Annotea also specifies a protocol for interacting with annotation servers to allow for the creation, maintenance, and discovery of annotation resources. In the context of the proposed project, both the creation and maintenance of annotations are considered out of scope, but discovery is clearly of interest. The query service proposed by Annotea does not specify a query language, and only supports the ability to request all annotations with a specified resource as the annotation target from an annotation server. Other functionalities as described in Batch Read and Selective Read, above, are not specified. Such capabilities are deemed important in the context of the proposed project. For example, Hunter’s HarvANA system described in a JCDL 2008 paper uses Annotea to model annotations but adds a Batch Read functionality based on OAI-PMH to aggregate annotations from multiple servers¹³. Hunter et al have also extended the Annotea model to support security, annotations of multimedia content and SPARQL querying to enable users to search and retrieve annotations based on author, date, type and content²². In addition, the proposed project intends to improve on the lack of adoption that Annotea has suffered through the work in the other project threads. These will anticipate support for the Annotation Model and Read Annotation Interfaces in existing, widely used tools and collections, rather than trying to enforce new behavior for the everyday tasks of scholarship.

A recent effort aimed at an interoperable scholarly annotation environment approach is that of the Huygens Instituut in The Netherlands. Their SANE project, Scholarly ANnotation Exchange¹⁴, created an XML Schema to express annotations, and a protocol for requesting

11 Agosti, M. Ferro, N., “A Formal Model of Annotations of Digital Content” ACM Transactions on Information Systems, Vol 26, Issue 1, 2007

12 Kahan, J., Koivunen, M., “Annotea: An Open RDF Infrastructure for Shared Web Annotations” *Procs of the 10th International conference on the World Wide Web*, pp 623-632, 2001

13 Hunter, J., Khan, I., Gerber, A., “HarvANA – Harvesting Community Tags to Enrich Collection Metadata”, *Procs of the 8th Joint Conference on Digital Libraries*, pp 147-157, 2008

14 <http://www.huygensinstituut.knaw.nl/projects/sane/>

them from an annotation server. The approach is clearly not Web-centric or resource-centric, and the lack of an XML-independent annotation model yields problems regarding adoption beyond the initial target group. SANE's focus on scholarly editorial practice has also biased the choice of XML elements (e.g. the 'selectedText' element), and the non-standard approach to address chunks of text (based on start and end points within the text). SANE's proposed protocol is a rather inelegant extension of OAI-PMH that reaches deep into the realm of searching, and fails to use a REST approach in cases where it clearly could. Still, SANE has clearly invested significant effort to consider the use of annotations for scholarly editions and the proposed project will aim at supporting the required functionality in its proposed model and information architecture.

Another noteworthy recent effort is the OATS – The Open Annotation and Tagging System – that aims at Web-based, collaborative tagging for educational purposes¹⁵. Although the OATS annotation model is not clearly documented, its REST-based approach for the implementation of some Selective Read functionalities (e.g. request all annotations by a certain creator) is inspiring. We also distinguish larger systems which include annotation capabilities, such as the general digital library models specified in DILIGENT¹⁶ and by Fox's 5S¹⁷ architecture, or attached as additional services to existing digital library models, such as the DiLAS¹⁸ service built by Agosti et al. and integrated with DAFFODIL and BRICKS. These models and services must especially be taken into account with respect to interoperability, as they form potential developers and early adopters, as well as sources of large quantities of annotations.

General Web annotation clients are often quite removed from scholarly practice but still very important to analyze for functionality and model requirements, as they have a large number of users and are the everyday tools to which many people are accustomed. There are a large number of such clients, with various functionalities, user interfaces, data models, architectures and target media formats. As the clients are not constrained by any requirement for interoperability or formal processing, they provide a very broad spectrum of how people envision the use of annotations online. Each must be analyzed for its intrinsic model, if discoverable, and it would be a very positive outcome if some or all of them were to adopt the model and/or architecture designed by the project. Notable examples of clients in this category include mystickies.com, diigo.com, sharedcopy.com and jotcloud.com.

Video in particular has seen a recent surge in annotation clients, including the well known YouTube annotation service, VoiceThread¹⁹, Project Pad from Northwestern University²⁰, Viddler²¹, and Hunter's Vannotea²². As more and more people annotate video, the rewards for having interoperable annotations attached will become more and more apparent, particularly with respect to searching and segmentation of video streams based on easily processable and time based annotations. This explosion of video annotation reinforces the requirements for

15 Bateman, S., Farzan, R., Brusilovsky, P., McCalla, G. "OATS: The Open Annotation and Tagging System". <http://fox.usask.ca/files/oats-lornet.pdf>

16 <https://twiki.cern.ch/twiki/bin/view/DILIGENT/DiligentArchitecture>

17 Goncalves, M., Fox, E., Watson, L., Kipp, N., "Streams, structures, spaces, scenarios, societies (5S): A Formal Model for Digital Libraries", ACM Transactions on Information Systems, Vol. 22, Issue 2, pp. 270-312, 2004

18 Agosti et al. "DiLAS: A Digital Library Annotation Service" International Workshop on Annotation for Collaboration, pp 91-101, 2005

19 <http://www.voicethread.com/>

20 <http://dewey.at.northwestern.edu/ppad2/>

21 <http://www.viddler.com/>

22 <http://www.itee.uq.edu.au/~eresearch/projects/vannotea/index.html>

seamless annotation interoperability between different media types, not just between text and images.

In summary, despite the vast body of work regarding annotation practice, annotation models, and annotation systems, little attention has been paid to interoperable annotation environments. The few efforts in this realm to date have either not been designed as Web-centric and resource-centric, or have modeling shortcomings that prevent any existing resource from being the source or target of an annotation, from any resource annotating multiple resources at the same time, and from giving an annotation an independent status as a resource itself. We feel that our past work in the realm of information interoperability (OAI-PMH, OAI-ORE, OpenURL, SRU/W) and the project team's expertise in the realm of annotation systems and content collections places us remarkably well to try and tackle the interoperable annotation problem.

IV. Statement of Work

The activities of this project divide naturally into three separate, intercommunicating threads:

- Thread 1: Annotation interoperability data model and specification
- Thread 2: Annotation application implementation and integration experience
- Thread 3: Scholarly annotation application analysis and model evaluation

Each thread will be assigned specific resources and work packages to complete. During Phase 1 of this project the dominant thread will be Thread 1 -- definition of data model and interoperability specification. The bulk of this work will be done at the Los Alamos National Laboratory Research Library; however, all of the other collaboration partners will support this work. Allen Renear at the University of Illinois at Urbana-Champaign and Jane Hunter at the University of Queensland (Thread 3) will lead annotation research at their institutions focused on the analysis of current systems and architectures and the understanding of annotation data models more generally. This work will inform the writing of the specification at Los Alamos. In parallel with this effort and preparatory for subsequent implementation phase of work (separate funding), staff at GMU and MITH will integrate AXE libraries into the Zotero application (Thread 2). As a primary target of adoption and deployment, timely development of the new AXE-Zotero annotation application is critical. This work will inform the definition of our shared data model and interoperability specification and will allow us to vet the potential adaptability of our specification for a primary target of opportunity.

In the final stages of phase 1 of this project, staff at all partner institutions will collaborate with staff at Los Alamos to develop test code and the preliminary components of a minimal reference implementation as a way to confirm correctness of data model and interoperability specification. Cross-walks and other mapping between shared data model and application-specific data models will pave the way towards future work to align real-world applications with the shared data model and interoperability specification and vice-versa. This preparation and vetting of the data model and specification prior to public release will ensure that the public release alpha specification and data model is compatible with architectures in use at partner institutions, and will provide an essential criterion for the decision to fund subsequent phases of work.

To ensure appropriate inter-thread communication, all co-PIs will convene virtually on a routine basis, by conference call at least once a month, and in person at least once during phase 1 of the project.

In order to obtain input from the community at large regarding the evolving interoperability specifications (in addition to feedback from the other project threads), Thread 1 of the project will follow the approach which was successfully used in the OAI-PMH and OAI-ORE efforts, consisting of the creation of a core group of voluntary technical experts, and open discussion on the Web. To a large extent, communication with this technical core group will be conducted remotely, but one in-person meeting of this expert group will be convened during phase 1 to allow for the development of broad consensus on the directions to be taken. This in-person meeting will coincide with the in-person co-PI meeting.

In order to keep the overall project direction synchronized with other evolutions in the realm of interoperability and scholarly communication, an Advisory Board will be created. This Advisory Board will be informed about project directions, and consulted as needed. At this time we do not plan on an in-person meeting of the Advisory Board, but there likely will be an overlap between Advisory Board members and the core group of technical experts to be recruited.

An overview of each thread, the work to be done by each thread, and each thread's budget, individual timeline and staffing is detailed in the following sections.

Thread 1: Annotation Interoperability Specification

Overview

This project thread focuses on facilitating the emergence of an interoperable annotation environment that allows leveraging annotations across the boundaries of annotation clients, annotation servers, and content collections. To that end, interoperability specifications will be developed in a process that involves community participation, and iterative vetting of evolving specifications in conjunction with the other project threads. The interoperability specifications will be made openly available on the Web under an appropriate Creative Commons license, and will detail:

- An **Annotation Data Model** for expressing annotations in an interoperable manner. This model can be adopted by annotation clients, servers, aggregators and other applications to create and share value-added annotations with all other conforming systems.
- An **Annotation Information Architecture** to define the requirements for and relationships between software components, including annotation clients, servers, aggregators and other applications.
- **Read Annotation Interfaces** to allow read access to annotations expressed according to the Annotation Data Model between components of the Annotation Information Architecture. These interfaces can be adopted by annotation servers to expose annotations to annotation aggregators; by annotation aggregators to further share value-added annotations with other aggregators and applications; by annotation clients to read from annotation servers, annotation aggregators and applications.

In order to guarantee an optimal integration of the interoperable annotation environment with both the Web and the Data Web, the specifications will be based on the primitives of the Architecture of the World Wide Web (Resource, URI, Representation), and principles from the Semantic Web, and the Linked Data effort. Wherever possible, and appropriate existing standards (e.g. W3C, IETF, ISO, OASIS, OAI) will be leveraged.

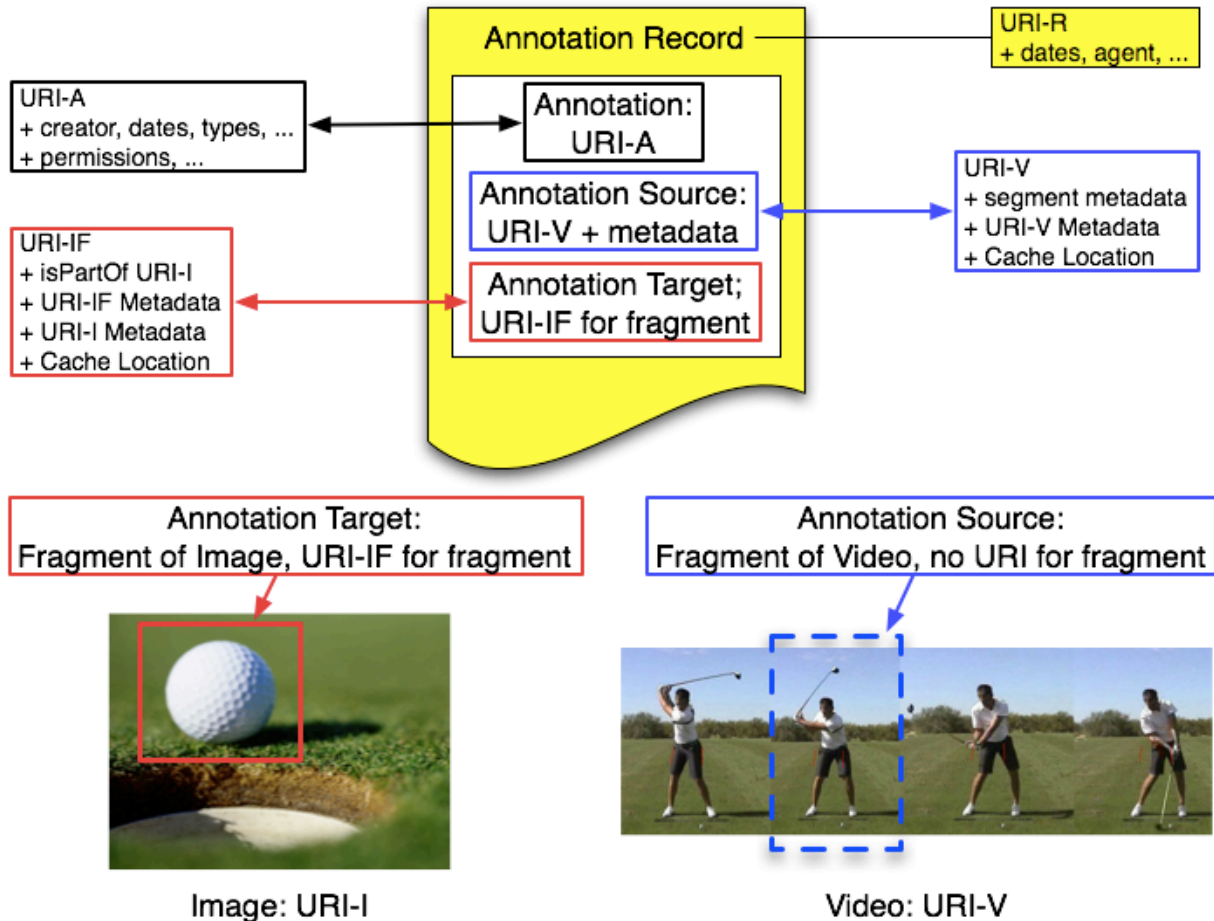


Figure 1: An initial perspective on a Resource-centric annotation model

This Web and Resource centric perspective on annotation interoperability yields the following initial perspective on the Annotation Data Model (Figure 1):

- Each Annotation is considered an autonomous Resource that groups a Source Resource and one or more Target Resources by means of an annotation relationship.
- The author of the Annotation, of the Source and of the Targets may be different. The Source may exist before the annotation relationship is created, or may be created at the same time as the relationship.
- Resources of any media type can be Sources or Targets. The Annotation is a conceptual Resource and has no media type.
- Annotations, as Resources with URIs, can be annotated with further Annotations.

This perspective is inspired by the interoperability direction taken by the recent OAI-ORE specifications, which were developed with support from the Andrew W. Mellon Foundation, and suggests that the Aggregation concept of OAI-ORE could be used as a foundation for the Annotation Data Model.

Figure 1 also indicates certain properties of particular interest that the Annotation Data Model will support:

- **Creator of an Annotation:** In addition to conveying the creator of an Annotation by means of a common name, the model will support expressing the creator's identity and the associated identity verification authority (e.g. a cross-organizational authentication environment such as Shibboleth or OpenID) under which the Annotation was created. This will support the emergence of trust-oriented services for aggregated Annotations.
- **Annotation Usage Permissions:** The model will allow expressing permissions pertaining to an Annotation, and these permissions may be leveraged by the Read Annotation Interfaces to selectively respond to requests depending on the nature of the client.
- **Cache Location of Source and Target:** Both URIs and the Representations available from those URIs are ephemeral on the Web. In order to facilitate maintaining the actual content of the Source and Target that was involved in a specific Annotation at its time of creation, the model will support expressing the URI of a cache location (for example in the Internet Archive) where that content is being preserved.
- **Segment Metadata for Source and Target:** In order to support Annotations that pertain to information at a granularity that is finer than which can be expressed by the URIs, media-specific fragment metadata may be required to support more specific addressing. Wherever possible, this will be avoided through the use of standards or well-known conventions for URIs that identify media fragments.

The anticipated interoperability environment focuses on sharing annotations in order to enable the creation of novel services based on annotations aggregated from multiple systems and environments. Therefore, the focus of the interoperability specifications will be on Read Annotation Interfaces; write, update or delete interfaces that are considered a concern of individual annotation client/server solutions. Read Annotation Interfaces will respond with Annotation Records expressed according to the Annotation Data Model, and two types of interfaces are anticipated:

- **Batch Read:** This interface will allow recurrently collecting batches of Annotation Records, using their creation/modification datetime as a selection criterion. Technologies such as OAI-PMH and Atom are possible foundations for this interface.
- **Selective Read:** This interface will allow collecting Annotation Records using other criteria that are relevant for the generation of services on aggregate annotation collections. Examples include requesting all Annotation Records that have a specific URI as Source or Target, that have a specific creator, etc. This interface will be REST-oriented, and technologies such as OpenSearch, SRU and OpenURL come to mind as possible foundations.

Figure 2 shows a possible Annotation Information Architecture with annotation clients interacting with their dedicated annotation servers using proprietary read-write-update-delete conventions. The figure also shows the introduction of an interoperable Read Annotation Interface adopted by several annotation servers. This allows the emergence of an annotation aggregator that collects annotations from compliant annotation servers, and creates value-added services with the aggregated information. It also allows the annotation aggregator to re-expose the original or

value-added annotations via another Read Annotation Interface. Now, annotation servers can use this latter Read Annotation Interface to collect aggregated annotations. These annotations can then trickle further down to annotation clients, using the Read Annotation Interface exposed by the annotation server. Following the *separation of concern* principle, these interfaces will not natively support selectively responding to requests on the basis of client permissions. Rather, an external technical layer, such as a cross-organizational authentication/authorization framework, will be leveraged to determine a client's permissions and the interface will be able to respond accordingly.




Figure 2: Annotation Records are collected from Read Annotation Interfaces

Thread 2: Implementation and Integration Experience

Overview

During phase 1 of the project, the Maryland Institute for Technology in the Humanities (MITH) at the University of Maryland, and the Center for History and New Media (CHNM) at George Mason University will undertake a thorough code review and restructuring of the Ajax XML Encoder (AXE), developed at MITH over the course of a one-year NEH Digital Humanities startup grant (2007-2008). These teams will also do a full analysis of the multi-year experience of Zotero users with Zotero's existing annotation tool to evaluate how the tool has or has not met expectations and use cases.

AXE allows users with limited technical knowledge to add metadata to text, image, video, and audio files located on the web. With an intuitive, web-based interface, AXE makes the process of linking digital artifacts more efficient and accurate. AXE is currently available as a working prototype on MITH's servers. However, late in the course of the development, it became clear that AXE should exist, not only as a stand-alone tool, but also as a code library that could be incorporated into other projects and tools (for example, as a plugin for WordPress). While this use is possible even in the current version, the code is not well documented, and because of the short timeline and relatively small amount of original NEH funding (about \$20,000 for one year after overhead costs), it was written with a goal of immediate functionality rather than generic reusability.

In the early fall of 2008, the AXE team spoke with the Center for History and New Media at George Mason University about the possibility of incorporating the AXE software library into the Zotero citation tool. The Zotero team was undergoing a reassessment of their in-house annotation tool with an eye toward the annotation of objects beyond web pages and toward the interoperability of Zotero annotations with the annotations created by other digital tools. Soon after the initial CHNM/MITH discussion, Chris Mackie and Ira Fuchs introduced the teams to Herbert Van de Sompel and the rest of the collaborators on this current project. It was clear that the AXE software library, with its ability to deeply link video, audio, text, and images located at remote URIs would be a useful test tool for the annotation schema the team planned to develop, and Zotero's immense user base would prove invaluable for both evaluation and promotion of the schema.

Thread 3: Scholarly Annotation Application Analysis & Model Evaluation

Overview

This project thread:

- Will participate in environmental scan of current annotation applications and repository service contexts, including defining cross-walk mappings between both the Pliny and Co-Annotea internal data models and the shared interoperable data model, plus identifying dependencies between MONK data store repository model and shared services needed to support interoperable annotation specification.
- Will support definition and writing of the shared annotation data model and interoperability specification.
- Will evaluate the model in terms of its ability to satisfy the current and anticipated future requirements of communities carrying out scholarly annotation using a set of different annotation tools

In accomplishing these objectives we will examine relevant research issues, address practical implementation issues, and lay the ground work for subsequent phases to adapt existing annotation applications and optimize existing content repositories to exploit the shared annotation data model and interoperability specification. Conducting in-depth analysis of implications for the MONK data store and the AusLit knowledge-base will allow us to examine potential implementation of the data model and specification over very granular and finely marked-up text, thereby anticipating the potential power of the model and specification and revealing implications of our modeling and specification work for repositories holding such content. Analysis of relationships between the Pliny and Co-Annotea internal data models and emerging shared, interoperable data model of annotation will allow us to examine potential of the shared data model to interact with such advanced and sophisticated annotation applications.

Research Issues in Scholarly Annotation Practice & Application

Annotation Service Functionality

We need to identify (beginning with an examination of current and emerging annotation tools) features of annotation that provide immediate, substantial value to both producers and consumers of annotations and annotated resources -- and then show our approach can potentially support these most essential features. Two issues are immediately apparent:

i). Annotation retrieval and filtering. This is a well-known objective for any treatment of annotations as "first class" objects. A scholar reading a text wants to retrieve the annotation about some portion of the text, and obviously they will also want to filter these by topic, date, annotator, etc. But which features are most important and must be foregrounded in our work? The general data model level will include an expressive and entirely generic attribute/value structure of course, but it is critical that a specific application be developed and encouraged. Moreover, the likelihood of uncontrolled vocabularies exploited by a variety of retrieval strategies must be anticipated and supported as well.

ii) The annotation paradox: Idealized Memory vs. Actual Value. We know that the actual exploitation of annotations by the original annotators fall well below their author's expectations for future use²³. How do we respond to this? Part of the answer may be improved functionality (finding, filtering, accessing, etc). If such improved functionality changes scholarly behavior that would be an enormous benefit. But we also may need to focus on kinds of annotation that do not fall prey to this paradox: annotations intended to be shared with others, which provide information entirely new to their consumers rather than being merely aids to memory and attention, or that serve a note taking function for scholars (see below).

Interoperability Data Model

The data model for interoperable annotation is of primary importance to the project, but what features are required to support scholarship? As well as the usual issues of completeness, expressive power, computational complexity, etc. there are additional problems:

i) Managing generality. Although our intuitive concept of annotation has a number of familiar prototypes, it is poorly delimited conceptually. Unconstrained general accounts can therefore result in data models that are simply instances of generic predication, designed to accommodate all intellectual activity whatsoever. If tool-building and adoption are to be successful, the focus must be on agreed upon refinements, and reference implementations, that match the most important and likely applications in the domain(s) of interest. Can classification of annotations²⁴ be of use in this context?

ii) Representing Reference (vs Anchoring). The question of what is being annotated by an annotation is distinct from how that annotation is anchored. This difference is poorly understood, and yet vital to the functionality of annotation sharing. For instance, an annotation may be anchored by layout features of a particular edition (FRBR: manifestation²⁵). But it is commonly features of the text (FRBR expression), such as words, phrases, sentences, speeches, quotations paragraphs, stanzas, etc. that are the focus of the annotator. At the same time whether the annotation is genuinely about the expression level features (the sentence) rather than a work level feature (what is said by the sentence) varies routinely: some remarks are philological, others philosophical, and some both. The consequences of the distinction are profound for shared annotation systems which must negotiate varying editions, translations, etc. Building on prior work at the University of Illinois and on models of citation such as implemented in Harvard's Canonical Text Services URNs²⁶, we will look for generalizable solutions consistent with our Web and Resource centric data model.

Scholarly Annotation Practices

What do scholars want to do involving annotation? It is essential to tie project outcomes back to communities of scholarly practice. There is a considerable body of work on relevant scholarly research practices which should be brought to bear on the development of any annotation

²³ Marshall, C.C. and Bly, S. 2005. Saving and Using Encountered Information: Implications for Electronic Periodicals. *Proceedings of CHI'05, Portland, OR, April 2-7, 2005*: pp. 111-120.

²⁴ Renear, Allen H., DeRose, Steve J., Mylonas, Elli, and van Dam, Andries. 1999. *An Outline for a Functional Taxonomy of Annotation*, presented to Microsoft Research, April, 1999. Available: <http://hdl.handle.net/2142/9098>

²⁵ IFLA Study Group. 1998. *Functional Requirements for Bibliographic Records*. International Federation of Library Associations / K. G. Saur: München. Available: <http://www.ifla.org/VII/s13/frbr/frbr.pdf>

²⁶ <http://chs75.harvard.edu/projects/diginc/techpub/cts-overview>

system for scholarly content²⁷. We are particularly interested in two key scholarly activities that require very different types of annotation—note-taking and netchaining—and that support very different kinds of work with content over time. To ensure widest adoption some effort must be devoted during the course of development to confirm our understanding of the specific domain practices (scholars using JSTOR) and to provide ongoing feedback with respect to the features of the system being developed.

i) Notetaking. One current system, Pliny (John Bradley, Kings College London), has made important advances on supporting note-taking and tagging activities performed by individual scholars. We aim to examine how this work might be extended to promote annotation sharing within a scholarly community studying the same primary materials. This kind of annotation functionality would support a process of intellectual accretion over time within an intellectual community or within a particular thematic research collection built for a scholarly community²⁸.

ii) Netchaining. In addition, we know that scholars perform “netchaining” across sources on the web, often with the aim of building their personal research collections of digital content. The IMLS DCC environment we've created is starting to support this practice to some degree with the addition of metasearch capabilities that allow users to move from primary sources, such as a photo of an event, to journal articles about that event. Annotation functionality is needed to support the intellectual work of keeping track of the value and relationships among materials retrieved, as well as the ideas generated. In particular, we are interested in what kinds of linking relationships scholars want to express about journal articles, or sets of retrieved journal articles, in relation to primary source objects or collections of objects, and how these connections can be represented. This kind of annotation is more likely to support a process of elimination or assessment over time.

²⁷ e.g., Carole L. Palmer, Lauren C. Tefteau and Carrie M. Pirmann. 2009. *Scholarly Information Practices in the Online Environment: Themes from the Literature and Implications for Library Service Development*. OCLC, online: <http://www.oclc.org/programs/publications/reports/2009-02.pdf>

²⁸ Palmer, Carole L., Cragin, Melissa H., and Hogan, Timothy P. 2004. Information at the Intersections of Discovery: Case Studies in Neuroscience. *Proceedings of the American Society for Information Science and Technology annual meeting* 41: 448-455.

