

UCLA

UCLA Previously Published Works

Title

Fluid Ontologies for Digital Museums

Permalink

<https://escholarship.org/uc/item/33j8k61m>

Journal

International Journal on Digital Libraries, 5(3)

Authors

Srinivasan, Ramesh

Huang, Jeffrey

Publication Date

2005

Peer reviewed

Fluid ontologies for digital museums

Ramesh Srinivasan, Jeffrey Huang

Harvard University, Graduate School of Design, 48 Quincy Street, Cambridge, Massachusetts, 02138, USA

e-mail: {rsriniva, jhuang}@gsd.harvard.edu

Keywords: Ontologies, Metaviews, Dynamic data structure, Digital museum, Online exhibits, Cultural heritage, Knowledge representation, Information visualization

Abstract

With the advent and accessibility of the Internet, artistic and indigenous communities are beginning to realize how digital technologies can be used as a means for documenting and preserving their histories and cultures. However, it is not yet clear what knowledge architectures are most appropriate for creating a digital museum in order to facilitate an effective collection, organization, conservation, and experience of cultural and artistic heritage. In this paper, we discuss the concept of “fluid ontologies,” a novel, dynamic structure for organizing and browsing knowledge in a digital museum. Fluid ontologies are flexible knowledge structures that evolve and adapt to communities’ interest based on contextual information articulated by human contributors, curators, and viewers, as well as artificial bots that are able to track interaction histories and infer relationships among knowledge pieces and preferences of viewers. Fluid ontologies allow for a tighter coupling between communities’ *interests* and the browsing *structure* of a digital museum. We present the key ideas behind the use of fluid ontologies within the context of digital museum design and seminal work in metadata/dynamic ontologies, particularly as it pertains to objects of cultural heritage, and discuss these characteristics in three concrete examples: (1) *Village Voice*, an online agora that ties together the narratives created by a group of Somali refugees using an iteration of community-designed ontologies, (2) *Eventspace*, a node-based collaborative archive for design activities, and (3) *Tribal Peace*, an online digital museum still under construction and evaluation that uses proactive agents to tie distributed Kumeyaay, Luiseno, and Cupeno reservations together in their quest to achieve greater political sovereignty.

1 Introduction

One of the key challenges in creating a digital museum is the design of its underlying architecture [1]. What should a museum collect and how should it organize cultural heritage? This is a problem that digital museums have in common with digital libraries, and much of the recent research efforts in both fields have focused on data models and metamodels for structuring digital informa-

tion [2]. Approaches that begin to consider the digital object as present within a landscape of other pieces of information, each maintaining its own attributes and interrelationships, have become critical. Parallels between the work of cultural heritage in the electronic space and the linguistic sign have become undeniable [3]. Indeed, recent years have seen a proliferation of powerful museum archives that are now more available than ever via the Web. Projects such as the Ellis Island Archive (<http://www.ellisland.org>) or Experience Music (<http://www.emplive.com>) have surfaced to tie past and present across a wide variety of individuals, be they immigrants or musicians. Concurrently, a significant amount of research is being done in the fields of knowledge representation and ontologies. These span deeply formal models of describing the world and the relationships among its attributes. For example, CYC, developed by Cycorp Corp., is one of the largest knowledge representation databases. Its goal is to construct the foundations of basic common sense to relate all the pieces of knowledge within the system. This involves a massive mélange of terms, rules, and relationships [4]. Another relevant project is the Open Mind Common Sense project (<http://commonsense.media.mit.edu/cgi-bin/search.cgi>), which is an attempt to provide computers with the millions of pieces of ordinary knowledge that humans know as “common sense.” Marvin Minsky, in *Society of Mind*, defines common sense as “actually more complex than many of the intellectual accomplishments that attract more attention...because the mental skills we call ‘expertise’ often engage large amounts of knowledge but usually employ only a few types of representations. In contrast, common sense involves many kinds of representations and thus requires a larger range of different skills” [5].

While both of these poles of projects are fascinating, we find a growing rift between the advances of AI-driven knowledge representation modeling and the increased digital presence of museums. Thus, we believe that a synthesis of these two poles holds great potential. In this vein, only a few projects have surfaced as useful benchmarks for our discussion. For example, The Dublin Core metadata project (<http://www.dublincore.org/projects/>) has attempted on a large scale to create a new interoperable set of standards to tie distributed content together by worldwide creators. The successes of this group are truly admirable and serve as a first solution in the quest to create cross-referenceable, deeper repositories of knowledge. This approach would enable curators to effectively create exhibitions of different presentation styles, with an understanding of how to select a variety of different pieces by treating them as communicative objects. Indeed, other standards such as IBM’s Digital Library (<http://www.surfnet.nl/innovatie/surfworks/doc/mmmetadata/h3.html#3.7>) and the W3C’s Resource Description Framework (<http://www.w3.org/RDF/>) exist to effectively integrate heterogenous content.

To explore existing metadata standards in further

detail, we point to and comment upon the following projects:

RDF/XML and potential for interoperability

The prevailing standards in terms of metadata schemas are XML and RDF schemas. The development of a mechanism of interoperability between the largely semantic, descriptive RDF and structural, datatyping XML would be very helpful. It is within this discussion that MetaNet is suggested to merge the domain-specific ontologies of RDF schema and XSLT is suggested to integrate the application profile instantiations of XML Schema. The development of a communication between these two merging systems could thus provide a much richer ontology to correspond to a given piece of information or object.

CIDOC/CRM

The CIDOC metadata model is focused on the integration and interchange of media within the diverse setting of objects of cultural heritage. As opposed to broader standards (RDF, XML), its ontological approach involves a restriction to the underlying semantics of database schemata and document structures within objects of cultural heritage. Thus, CIDOC/CRM is a formal model to integrate, easily parse, and serve as a common language/concept for museums that wish to maintain heterogeneous digital collections (<http://cidoc.ics.forth.gr/>).

ABC ontology and model

Funded by the Harmony project, the ABC Model was developed with the express purpose of providing interoperability between different metadata ontologies and enabling communities to begin to develop their own descriptive ontologies. It has become the emerging standard for resolving some of the issues raised in the RDF/XML discussion and uses formal logical categories such as situation, temporality, and object-oriented relationships [6].

Other models

Other more formal models include (1) OWL Web Ontology Language (<http://www.w3.org/TR/owl-ref/>), which is a vocabulary extension of RDF used for publishing and sharing ontologies across the World Wide Web; (2) XML Topic Maps (<http://www.topicmaps.org/xtm/1.0/>), which provide a model and grammar to represent the structure of information resources used to define topics and their associations. This is basically an XML-customized, interoperable, relational ontology language.

Where we stand

The approach presented in this paper offers a shift from these research efforts in that instead of focusing on the development and definition of clear but fixed standards for structuring information, the main motivation for this research has been in liquefying such structures and designing “fluid ontologies,” i.e., processes for letting knowledge structures emerge from the interaction with the very communities that are using the digital museum. We believe that an ontology that is truly adaptive and reflective of the priorities and hierarchies of the participants (museum visitor, curator, or contributor) can serve as the key architecture behind making the digital museum experience truly powerful. Rather than subscribing to the more formal models discussed, we focus on different, more anecdotal (as in the case of Village Voice), self-created (as in the case of the metaviews of EventSpace), and quantitative mapping approaches (as in the case of the proactive agents used within Tribal Peace). In all these cases, the question of sacred or private content was carefully considered. Community members would submit pieces with the understanding that their submission would be shared only among those who had login and password information (restricted to community members participating in the project). In the case of Village Voice, yet another version of the system was created to serve as a completely open and public portal where any Internet visitor could access the system. Community members comfortable with sharing content on this level were enabled to submit pieces for this version as well.

In the remainder of this paper we describe the key ideas behind fluid ontologies, show how these ideas are realized in sample online cultural heritage and artistic community projects, and discuss their implications for the future design of digital museums.

2 Key ideas

One of the most important characteristics of fluid ontologies is that they are not predefined but emergent (and adaptive) structures for knowledge representation. Four key ideas form the basis for fluid ontologies. Although these ideas can be seen as independently and empirically testable hypotheses, they are treated here as general design premises. They are listed and briefly described below. In the next section, we describe concrete projects that use them.

1. Involvement of content creators: The ontology can become richer if the content creator is directly involved in its definition. By letting the creator or collector of content participate in the definition of the knowledge structure and indicating where his or her piece of knowledge fits in, we gain additional insight about the context for the new content and how the pieces relate to each other. Individual

pieces of knowledge no longer stand alone in this model but instead are interrelated in the collective fabric of the entire exhibit.

2. Metaview sharing: The browsing and learning experience can be significantly enriched if the views each participant creates of the available knowledge and the way the participant makes sense of, browses, and rearranges the knowledge are made explicit visually and can be accessed by others.
3. Adaptiveness: Ontologies that adapt and are continually redesigned in time can be more useful than static ontologies. How people make sense of the world often evolves over time as a function of what they experience.
4. Bots and personalization: The use of artificial bots and proactive agents that track and analyze interaction histories can effectively help adapt and fine-tune the dynamic evolution of ontologies. By integrating and learning from interaction histories, we can dynamically generate ontologies corresponding to the different viewers' profiles and preferences and anticipate expectations and interests.

3 A first example: Village Voice

To illustrate the ideas of fluid ontologies described above, we provide in this and the two next sections examples of how these ideas have been applied to real projects for the collection, preservation, and exhibition of cultural heritage. The first project is Village Voice, a dynamic, expanding archive of community-submitted narratives developed as part of Srinivasan's M.S. research at the MIT Media Laboratory [7]. This project was designed, developed, and evaluated within a Somali refugee community based in the Boston metropolitan area.

3.1 The Village Voice ontology

Village Voice serves as a useful first example of fluid ontology because it is structured to represent its content based on the manners in which content creators (community members) articulate their own cultural realities. In addition, it builds upon the growing movement in community publishing and storytelling. A number of studies have demonstrated that empowering communities to create their own stories stimulates a process of reflection, which in turn facilitates the sharing of values, knowledge, structure, and dreams [8]. However, it also builds on the concurrent technical movement in knowledge representation and ontology-focused research. Village Voice, then, is the intersection of how to apply the power of a fluid ontology to a set of community-created narratives.

In the case of Village Voice, ontology can be seen as a conceptual map where the links between individual

pieces of knowledge are delineated. An assumption researchers in this field make is that knowledge is without meaning unless it is contextualized. The specific nodes in the structure need to be understood along with the links that tie them together. Roger Schank explains that through ontology we make sense of the world. Information that we encounter is understood through our own internal “data structures,” which he calls scripts [9]. In contrast to these ambitious projects that attempt to map universal knowledge, the Concept Maps approach involves the creation of diagrams that can represent specific concepts and their interconnectedness [10].

The idea of creating formal, relational links between concepts was an inspiration to Village Voice. Instead of basing community contributions and their interrelations around ad-hoc indices, Village Voice would allow the community member or other user to interact with material based on how the community itself articulates the relationships within its different pieces. When continuously populated with their stories, ontology becomes a dynamic structure that is used by members to model the evolution of their community. As the focus of the study was to test whether a community-created model could be a strong foundation for the design of such a multimedia narrative system, Srinivasan sought out the Somali refugee community in the Boston area, a population of approximately 3000 to 5000. Refugees today are victims of a civil war that has torn apart these families and decimated a once-thriving culture. This community has dramatically expanded over the last five years due to the civil war in Somalia. According to community members, there is a desire to archive their experiences as they face new challenges in the United States. They wish to find a means to tell stories to their community, as well as to incoming refugees and others outside of the community. Traditionally, story has been orally transmitted in Somali culture, so the use of a medium that records and retells story is new to them.

Thus, a set of video stories was created by a variety of community members. Very few instructions were given about the pieces except that they should be focused on issues relevant to their collective community. The goal was to use these stories to stimulate the design of a set of representations, or a fluid ontology, that could illustrate the intersecting issues of the community and could be altered and modified as the community’s priorities evolved.

As the community would meet over several sessions the role of the fluid ontology became clearer. Each session began with an explanation in Somali to remind participants about the purpose of the project. The movies were then shown on the VHS tape, and participants were encouraged to pause, stop, or repeat the video at any time. They were instructed to do this whenever they felt an issue that was relevant for their community was revealed in a story. During the pauses, the community

would discuss the videos they were watching and craft a part of the ontology diagram on the white board in the front of the classroom.

During these discussions, the community would come to a consensus on whether an issue that had come up should be included in the ontology. For example, one story was set at a Somali youth party. It showed teenage men and women dancing together dancing to hip hop music. The idea of a youth dance party without Somali music was disagreeable to some of the participants because of its disrespect to the Islamic taboo of premarital relationships, while most of the youth at the meeting argued that one could have a premarital relationship without being disrespectful to Muslim culture. During this discussion, the participants decided that issues of religious tradition, sexuality, and generational differences were relevant to the ontology. These topics were then added to the ontology and linked to each other on the white board. The drawn structure changed multiple times in the process, as the community members reflected further on the issues that united them. The initial community ontology appears as follows:

The video content was structured based on the ontology using a simple MySQL database and PHP server-side scripting. As is made clear, the relationships are specified as set – subset, and no deeper information on relationships between topics has been specified. Every term within a box in Fig. 1 was a concept (object or practice for example) relevant to the Somali community, and terms connected to the box as subsets were considered specific instantiations of the above box. This community ontology was generated through the concept mapping [11] approach, as described earlier. Village Voice's use of Concept Maps is based on the theory that this technique could also serve as a useful simple architecture to interrelate the narratives created by community members within the digital system. Thus, the implementation of ontology in this case was discrete. While the community could meet at any time to redesign their collective ontology and implement this within a simple PHP script, the ontology was not truly dynamic and able to sense changes within community priorities on a continuing basis. It is largely as a reaction to these shortcomings that the collaborative filtering agent described in Sect. 5 was developed.

Community members were thus able to submit their video pieces through another PHP upload page that would query them to select their video and representative thumbnail image and select a set of checkboxes, each of which represented a term within the community ontology from Fig. 1. Thus the video would be submitted along with a set of attributions it maintained to different concepts within the collective community ontology. The upload page appeared as follows.

It is important to note that this was merely a first iteration within the process of creating a set of fluid

ontologies to anchor the content that began to populate Village Voice over time. However, as we assert in Sect. 2, Village Voice incorporates several attributes of our characterization of fluid ontologies. Its explicit focus on author/community-articulated knowledge design enables the contributor to be involved in the exhibit design to an unprecedented degree. Additionally, the potential to share metaviews during this process of iterative ontology design is made possible through the process and technology discussed. Finally, it is important to state that within this project the community of authors is given the ability to continually redesign and update the collective ontology. As priorities and goals change, the collection of narratives are empowered to change their relative meanings. This creates a level of fluidity where the exhibition can be responsive to the iteratively updated knowledge architecture.

3.2 Prototype technology

To integrate the content submitted and fluid ontology, a simple Web portal was created.

Village Voice is entered via a “search” page. This page is organized hierarchically, according to the tree-based design of the fluid ontology. The user can select multiple nodes from the ontology tree that he or she is interested in watching stories about.

This leads to the main browsing interface, designed to illuminate complex interrelationships between different topics visually. The browsing page is designed to give the user a wide range of information about the different video stories by displaying each story’s thumbnails while conveying their relationships to each other. Each thumbnail can be selected, so that it can be streamed in an adjacent frame, or so that audio responses from other community members can be listened to. As seen in the image below, the thumbnails are illuminated to varying levels. This is reflective of how closely each thumbnail corresponds to the terms the user decided to search on in the search page. A brighter illumination indicates a closer match with the search query.

3.3 Evaluation

To assess whether the fluid ontology serves as a more attractive design for the visitors and contributors of the digital exhibition, a simple experiment was designed. Over a week, approximately 100 community members were each asked to browse the ontology-focused Village Voice and a control keyword version that would only organize stories based on the words most often stated within them [12]. Before using either version, the subject would log on to the system with an anonymous name allowing the ability to monitor the sequences of stories the subject browsed, how long he or she stayed logged on, and how many stories were played.

These data show the mean and standard deviation values across the subjects of the number of stories browsed, number of stories played, and time online for the keyword (KW) and Village Voice (VV) versions. The values for time online are expressed in terms of seconds.

These data show a higher engagement for subjects across the board with the fluid ontology version of Village Voice. It appears that fluid ontologies inspire a greater engagement with the digital exhibit. This example thus begins to confirm our assertions that fluid ontologies can serve as an empowering mechanism for the creators and visitors of the digital museum.

4 Example two: Eventspace

Where Village Voice explored the idea of a community-driven ontology that is designed and redesigned based on consensus, our next example, Eventspace, explores the ideas of fluid ontologies within the context of a bottom-up online exhibit. In this example, a multiplicity of ontologies coexist in parallel that are created in a decentralized fashion. The objective of using fluid ontologies in this fashion is to collect and exhibit not only the evolving artifacts created by the community of artists but also their processes of creation and perspectives. We want to enable individual artists to do three things:

1. Place their contributions in the larger context of a digital museum, i.e., precisely situate where their own contributions fit in;
2. Structure the “larger context” of the online exhibits from their own perspectives and in their own voice, i.e., create and communicate their “metaviews” of the existing artifacts to others; and
3. Use those views to stimulate collaboration on art works, i.e., based on an increased understanding of where each contributing artist is coming from, encourage joint projects.

4.1 System overview

System architecture. Eventspace evolved in two stages. The first version was written and refined between 1995 and 2002 using open-source software at the Swiss Federal Institute of Technology: MySQL, Apache Server, PHP, and HTML. This version is described in [13]. The second version, currently under development at Harvard Design School, Center for Design Informatics, uses a Flash front-end, Actionscript/Coldfusion, and Oracle database. It currently exists in a prototype form and has been tested in two academic experiments.

Functionalities. Eventspace features several functionalities that empower members of the online community

to post digital artifacts in an open online exhibit visible to all participants. Each submission, or node, consists of a particular set of required elements, as specified by the community leader or curator. For example, a digital arts exhibit might define a node as a combination of a graphics file (image, 3D model, animation), URL, and text description. The database automatically records for each node submitted its author, time of submission, and whether it was an original “creation” or a continuation of a previous node.

Eventspace has built-in uploading mechanisms that enable participants to contribute ongoing exhibits in one of three ways: by posting submissions (nodes), by building on the nodes of others, or by commenting on the submitted work. This system builds on the understanding that each participant brings different skills to the exhibit: creators, advancers, and reviewers. This is not unlike the roles typically found in design teams: some are best at producing original ideas (creators), others at building on the ideas of others (advancers), while still others are most skilled in critiquing the work submitted (reviewers). Every node retains a record of (a) all the previous nodes from which it is derived, (b) all the comments made concerning that particular node, and (c) all the later nodes that built on ideas or files contained in the current node.

Default ontologies. What emerges from such a structure is a chronology of the emerging online exhibition: a visualization of the “genealogy of ideas.” This genealogy is a graphic representation that clearly identifies a network of individual submissions leading to the final collectively authored exhibit. Eventspace encourages interaction and collaboration among participants. It must be noted that such collaboration also contains an element of Darwinian competition: when artists can easily view and compare all the ideas (nodes) proposed by their peers, there will be a process of natural selection as participants choose which ideas to build on and which to neglect. Eventspace adds to any online exhibit the ability at any stage in the creation process to make completely transparent the authors who have contributed to a given submission and the ideas that have been most influential in shaping the outcome of the collaborative project. Participants can choose three basic representations of the nodes as defaults: organized by time (chronological genealogy of ideas), by people (who contributed what), and by keyword (as specified by contributors).

Metaviews. By the metaview function we denote the ability provided by Eventspace for each participant to create and recreate his or her view of the current set of knowledge nodes at any given time and share it with the community. In other words, in addition to the default static ontology with its three default views, Eventspace provides each participant with the ability to dynamically

construct a custom view and arrangement of the existing nodes, or what we call “metaview” of the existing body of knowledge, and share those views with the community. We believe that seeing and comparing multiple perspectives of the same set of knowledge adds tremendously to the experience of the online exhibit. Since everyone can create, edit, and add metaviews, multiple evolving ontologies are available at any given time for accessing the same set of knowledge. Technically, this is accomplished by making it easy for participants to script custom views, or so-called “templates”, onto which the content of the nodes are pushed, reflecting the database-driven architecture of the site.

4.2 Sample application

Our first and most fully developed application of these ideas is in the context of two class projects at Harvard Design School and the Department of Visual Arts that we use to test our ideas of fluid ontologies.

Background. We chose the academic context as a testbed because of its accessibility and controllability and because making work in progress explicit and learning from one’s peers have traditionally played a significant role in the education of artists and design professionals. Artistic and design knowledge is tacit by nature and difficult to make explicit: it cannot be lectured in a classroom and is best taught through hands-on experience, trial and error, and peer-to-peer interaction. The physical space of the artist’s or designer’s atelier was specifically designed to meet these pedagogical challenges of teaching. Its open studio space fosters the constant exchange of ideas around student works in progress that is an essential aspect of the arts studio experience. Eventspace in this context encourages peer-to-peer interaction by engaging participants in collective projects that demand the combined resources of several students working as a group. These groups can be assigned by the instructor, or they can be formed by the students as the need arises for collaboration on a particular task.

In our test cases, we used Eventspace in the following three ways, which we introduced gradually as the semester progressed.

1. First, we introduced Eventspace as a collective curation tool. Participants were asked to post nodes containing interesting art works that related to the subject of the class and post them as nodes on the online exhibit.
2. Second, we used Eventspace as an online conversation tool to stimulate discussion around the nodes posted. Here participants assumed different roles, that of creators, advancers, and reviewers, as mentioned above. Creators would continue to post new nodes that may contain their own arts projects,

advancers would take an existing posted node and develop it further, and reviewers would critically comment on an existing node.

3. Finally, we introduced the metaview idea. Each participant would create a metaview of how they viewed the nodes and nodal relationships. Such metaviews could be “snapshots” of the current exhibit, static views that show how individuals view the current set of nodes in one particular moment in time, or “dynamic filters”: templates for structuring and organizing changing current and future nodes.

4.3 Discussion

As of this writing, the described application has been in use for 4 months. Compared to the Village Voice application, the main difference here is the interpretation of fluid ontologies in terms of metaviews. By providing the capability to dynamically create, modify, and add new views of the same set of knowledge, multiple evolving ontologies simultaneously coexist. Participants thus have multiple fluid ontologies at their disposal that they can interchange to browse the online exhibit, including three default views provided by Eventspace – chronological, by author, by keyword – and the 20–25 metaviews created by participants. The following diagram illustrates the overall architecture of the Eventspace as deployed test cases:

5 Example three: Tribal Peace

Our final example will be considered merely a side note, given that it is an ongoing investigation into how fluid ontologies can begin to take on a more proactive role in the context of social, educational, cultural, and political empowerment. While the previous two examples have studied fluid ontologies strictly in terms of a set of iterations or metaviews, Tribal Peace utilizes these techniques as well as the approach of a proactive agent that can arm individuals across reservations with novel and useful information collected in a large digital museum archive. The ongoing project involves the Kumeyaay, Luiseno, Cahuilla, and Cupeno Native American reservations distributed across Southern California and Northern Baja California (Mexico). Waves of Spanish, Mexican, and, finally, American conquerors have had a strong effect on a number of tribes within the San Diego region of Southern California. The reservation systems of the 19th and 20th centuries have magnified these dynamics, forcing a resettlement of these groups away from their traditional coastal lifestyle into the desert approximately 100 miles to the east [14]. Specifically, the situation is now that the different tribes of southern California have largely

fragmented and maintain little infrastructure from which reorganization could be made possible.

In an effort to assist these communities and with a Hewlett-Packard grant, the Tribal Digital Village project was begun. This project has attempted to study whether the provision of basic satellite-generated Internet facilities could enable individuals across tribes to receive the informational benefits the Internet makes possible while also creating a digital museum that could empower the dispersed tribes to organize and communicate regarding their ongoing quest for a greater level of political sovereignty. This digital museum is continuing to receive tribal member-authored content, the only restriction around which is that it be reflective of the political, cultural, or educational issues that impact the communities in general. This can involve, for example, a peek into the day-to-day life of a certain tribe member. However, it can also be considered a documentation of an ancient tribal practice, such as the return to the sea pilgrimage. What is important is that different topics and issues that are important to the tribe be integrated into the collection process. Additionally, it is important that content be created by a cross-section of community members across traditional social boundaries such as age and gender.

However, despite the growing presence of this digital museum, there exists no architecture or mechanism to narrate this large aggregation of content into an experience that can empower different museum visitors and tribal members to receive meaningful exposures to content that can allow them to connect to others around the important struggle toward sovereignty around which a number of these dispersed reservations identify. This presents a situation that calls for an application of a fluid ontology that can serve as an architecture to present a meaningful experience to the visitors of the digital museum.

Traditional digital archives are merely repositories; they are not necessarily active information generators. Therefore, in this context, because tribal leaders such as Anthony Pico, chairman of the Viejas Band, are envisioning the Web portal as a resource in their fight for sovereignty, the archive must be much more proactive and connection-brokering. The approach being studied in this example involves the use of fluid ontologies in the form of proactive agents that can mediate and that can disseminate knowledge more powerfully among the visitors that access the digital museum. We define agents as technological entities that attempt to accomplish a task for a user based on an understanding of the user, the content of the task, and the environment in which the task is situated (other agents, the information being expressed, etc.).

Specifically, the focus is on an intelligent agent or bot capable of mapping the evolving interests of a given tribal individual to a meaningful and previously unseen set of

content that can be presented to him or her. This approach has been touched on briefly in Sect. 2 above in our discussion of how the approach of fluid ontologies makes possible digital museums that can learn from interaction histories and dynamically generate ontologies that correspond to the different viewers' profiles and preferences while also anticipating expectations and interests.

Our utilization of agents departs from the traditional research trajectory where the agent is only designed as a representation of an individual within a usually economic, numerical context. Instead, we are interested in how the agent can actually create a connection between a given community member and aspects of the quest for sovereignty that may appeal to him or her, thereby creating a structure around which the struggle for sovereignty can be structured, discussed, and disseminated. We have focused the creation of the proactive agent around several notable examples, including Oren et al.'s Guides project and Maes et al.'s work on social-information-filtering algorithms. We are in the process of assessing the relative value of this recommendation agent for Tribal Peace by studying time-based changes in usage patterns, levels of knowledge among tribal members, and social networks. The initial results of this work should be available in June 2005.

6 Benefits and potential problems

We have briefly described three applications of the fluid ontologies in completed and ongoing projects. Across the three examples the following observations can be made with regard to the perceived benefits of using fluid ontologies

a. Benefits to the visitor

- A first significant benefit to the visitor is the personalized interfaces and views of the repository, i.e., the ability for visitors to see content that they have not yet been exposed to but would find truly rewarding. This is particularly evident in the third example, where a bot makes recommendations based on its understanding of the individual preferences and interaction histories in the community (collaborative filtering).
- A second benefit derives from the new ability to access content and paths through content based on the articulation of the community itself; in other words, a visitor's ability to see content through the communities' own eyes (village voice and eventspace).
- A third benefit concerns the tighter coupling of social networks, i.e., the ability to be exposed to the community of creators and other visitors through the system's capacity to proactively identify connections with others.

b. Benefit to the creator/community of creators

- A first benefit to the author is the ability to get a sense of where his or her contribution to the exhibition fits within the overall structure.
- A second benefit as a community of creators is the ability to create an exhibit, not just content but the way the content is represented as an architecture, empowering potentials as the Village Voice evaluations have shown for a particular community of creators.

On the other hand, we also see potential problems associated with the use of fluid ontologies. Here are a few:

- The ontology is not truly fluid in the sense that it is a set of discrete steps rather than a truly continuous process like what a bot makes possible.
- The process by which the ontology is arrived at is imprecise.

7 Conclusions

In this paper, we have presented key ideas and sample applications related to fluid ontologies. We have shown how the careful design of the degree of “fluidity” of an ontology can provide a basis for powerful digital museum architectures. We believe that systems that use fluid ontologies, like those presented here, illustrate important – and not yet widely recognized – possibilities for creating, organizing, preserving, and exhibiting cultural and artistic knowledge. The power of the fluid ontology approach appears to be partly due to the fact that it does not emphasize the building of intelligent, a priori standardized knowledge structures but instead focuses on creating processes and providing tools to gradually support the sense-making processes of humans when they are confronted with cultural and artistic heritage.

References

- [1] Takahashi J, Kushida T, Hong J-k, Martin W, Rieger R, Gay G, Sugita S, Kurita Y, Reeve J, Loverance R (1998) Global digital museum: multimedia information access and creation on the Internet. In: Proceedings of ACM Digital Libraries '98, Pittsburgh, PA, 23–26 June 1998, pp 244–253
- [2] Trant J, Bearman D (2003) Foreword. In: Proceedings of Museums and the Web, Charlotte, NC
- [3] Smiraglia R
<http://phoenix.liunet.edu/%7Esmiragli/nature.html>

- [4] Guha R, Lenat D et al. (1990) Building large knowledge based systems. Addison-Wesley, Reading, MA
- [5] Minsky M (1986) The society of mind. Simon and Schuster, New York
- [6] Lagoze C, Hunter J (2001) The ABC ontology and model. *J Dig Inf* 2(2):Article No. 77, 2001-11-06
- [7] Srinivasan R (2000) Village Voice: expressing narrative through community designed ontologies. Massachusetts Institute of Technology, Department of Architecture, Program in Media Arts and Sciences
- [8] Campbell J (1988) The power of myth. Doubleday, New York
- [9] Schank R (1992) Agents in the story archive. Institute for Learning Sciences, Northwestern University, Evanston, IL
- [10] Novak J (1998) The theory underlying Concept Maps and how to construct them.
<http://cmap.coginst.uwf.edu/info/>
- [11] Novak J (1998) The theory underlying concept maps and how to construct them.
<http://cmap.coginst.uwf.edu/info/>
- [12] Salton G (1968) Automatic information organization and retrieval. McGraw-Hill, New York
- [13] Spicer DE, Huang J (2001) Of gurus and godfathers: learning design in the networked age. In: Education, Communication and Information, Routledge: Vol.1, No.3
- [14] Shipek F (1986) Pushed into the rocks. University of Nebraska Press, Lincoln
- [15] Oren T, Solomon G, Kreitman K, Don A (1990) In: Laurel B (ed) Guides: characterizing the Interface in the art of human computer interface design. Addison-Wesley, Reading, MA
- [16] Shardanand U, Maes P (1995) Social information filtering: algorithms for automating “word of mouth”. In: Proceedings of CHI 1995, Denver, CO, pp 210–217

Fig. 1. Somali community ontology

Fig. 2. Village Voice upload page

Fig. 3. Ontology-based browse page

Fig. 4. Browsing page: collage, video streaming, annotations, and more

Table 1. Mean and standard deviation values of Village Voice vs. control

	Mean value	Standard deviation
KW time online	263	205.36
VV time online	967	891.39
KW # clips browsed	2	0.697
VV # clips browsed	7	2.719
KW # clips played	1	0.433
VV # clips played	3	0.788

Fig. 5. A single submission, or node, showing (a) *left*: previous nodes on which this node was built and (b) *right*: later nodes built on the ideas or files implicit in this node

Fig. 6. Sample metaview by one of the participants (Mark Meagher): Visualization of the available knowledge node as a three-dimensional citylike visual structure. Each “building block” represents a knowledge entry. On mouse click the content is shown in the *right window*, together with related nodes

Fig. 7. Diagram of the Eventspace architecture in the test application

Fig. 8. A snapshot of the Tribal Peace system interface