THE DESIGN AND IMPLEMENTATION OF A REPOSITORY FOR DIGITAL CULTURAL ARTEFACTS

SUBMITTED: February 2004
REVISED: May 2004
PUBLISHED: July 2004 at http://www.itcon.org/2004/10/
EDITOR: B. Martens and W. Jabi

Robert Woodbury, Professor
School of Interactive Arts and Technology, Simon Fraser University
email:rw@sfu.ca

Michael Docherty, Director
Information Environments Program, School of Information Technology and Electrical Engineering, The University of Queensland
email:mjdoc@uq.edu.au

Hank Szeto
Information Environments Program, School of Information Technology and Electrical Engineering, The University of Queensland
email:hanks@espc.com.au

SUMMARY: *We motivate, describe and demonstrate a generic repository for visual material related to cultural disciplines such as architecture, art and heritage. The repository design abstracts from the requirements of particular tasks to implement a set of common features based on the metaphor of a gallery. In functional terms, galleries are spaces in which several roles play together to create the larger social entity we call a gallery. Curators organize and interpret collections. Exhibitors contribute and interpret collection material. Through their commentary, Critics relate Resources and Exhibition. Viewers visit galleries to read and make their own interpretations of the content therein. Our design provides for each of these roles and implements them as actions on five fundamental types of objects: Resources, Exhibitions, People, Works and Annotations. We describe the architecture of a system called A•VI•RE that supports the above four roles and five object types. Included in this design is a meta-metadata scheme that supports variation in descriptors across content types. A•VI•RE's implementation is based on the open-source web portal system TikiWiki, providing opportunities for customization and integration with other open-source efforts. We describe a set of initial authoring styles for A•VI•RE Exhibitions. The paper concludes with hypotheses about the social and intellectual opportunities provided by the A•VI•RE functionality.*

KEYWORDS: *gallery, digital media library, learning repository*

1. BACKGROUND

For several years we have developed and supported digital repositories of visual material for architectural libraries, scholarship, teaching and learning. What emerged as common from these projects is how the value of collections of visual material can be enhanced by user-provided commentary. Users act as authors, and their productions add value to collections.

Our past systems included early teaching websites; a teaching website compiler we called *WebWeaver* at The University of Adelaide, the *DigiLib* system at the University of Queensland (UQ, 2004) and the *vGallery* at The University of Adelaide (Shannon et al., 2001, Woodbury et al., 2000).

In particular, the *DigiLib* project aimed to enhance teaching and learning in visually based disciplines such as architecture through the creation of open access to architectural image collections and to embed the image access into the problem based learning of the discipline. Students and academic staff are able to access the image collection via the Internet and to select sets of images forming the basis of a lecture series or to be included in project and assignment activities.

vGallery is a web-based tool for aiding interactions amongst groups of learners (students and faculty) based around student work and critiques thereof. It supports faculty in creating virtual spaces for collaboration—taking the term 'collaboration' to cover a range of meanings from the private presentation of a piece of student work for formal assessment to the posting of work for general viewing and discussion (online or face to face).

A related project is the Digital Image Database (DID) at James Madison University (Sharon, et al., 2002, Pitt, et al., 2002). This is a free system developed by the School of Art and Art History, focusing on providing an instructional tool. As well as the obvious online content search library, it is also an effective online study tool for students and an in-class presentation application. Though our initial workpredated publication of the Madison DID, currently the DID represents a model of the usability and workflow appropriate for online images for visually based disciplines.

'Greenstone' (Witten 2001) is an example of open source DL software allowing users to build their own digital repository from existing resources. It comprises various modules for creating and maintaining a repository, and has an impressive list of users. Notably, it is supported by UNESCO, and represents an effective way of making digital libraries available to a widely disparate user group.

What the *DigiLib* project highlighted was the need to source accurate metadata, and emphasised the logistical problems in generating that data. This in turn led to the design, implementation and population of a generic system. Based on a gallery metaphor, the system design supports collecting, archiving and interpreting sets of digital, visually-based material.

2. OBJECTIVES

Our overall objectives are to devise, implement, use and test a generic repository for visual material related to cultural disciplines such as architecture, art and heritage. Our design abstracts from the particular requirements of, for instance, slide collections, lectures or student work-posting, and implements a set of common features based on the metaphor of a gallery.

In functional terms, a gallery can be conceived of as a *space* in which *curators* organize and interpret collections of material. The material may be sourced by a curator or may be contributed by an *exhibitor*. *Viewers* may visit a gallery to experience its content and to read the curatorial interpretation of the collections shown. *Critics* provide commentary on the material and on the curatorial interpretations of it. What was needed was a system design that provides for each of these roles and implements them as actions on five fundamental types of objects as follows. *Resources* are atomic objects in the system. Resources comprise any digital object (including ones distributed across logical files). *Exhibitions* comprise collections of Exhibitions and resources, along with authored interpretation of the Exhibition contents. A gallery can be thought of as a single Exhibition that contains the Exhibitions comprising the gallery's content. Most resources describe a *Work*, that is, an object in the world outside of the gallery. Works take a variety of forms. Buildings, un-built projects, books, paintings and sculptures are all examples of works. Works may also contain other Works as their proper parts. *People* access Exhibitions and both Resources and Exhibitions cite the People involved, be they, for example, authors, curators, creators of works, custodians or critics. Resources and Exhibitions may be tagged with *Annotations.* These function as lightweight commentary on a gallery's contents.

Each of these functions is described in more detail in the Section 4—User Scenarios.

3. METHODOLOGY

Building a gallery is both a technical and social enterprise. The technical issues include a metadata scheme, database design, interface design, implementation and meta-tagging process. Social issues are crucial and revolve around discovering and designing workable processes for meaningful contribution to a gallery system.

A metadata scheme must describe a gallery's contents. Currently textual metadata is the sole handle for information retrieval from a gallery. Our metadata scheme is developed from several sources, including the Visual Resources Association's (VRA, 2004) Core standard. Our metadata scheme is marked by its separation of *Works* from *Resources,* and by its inclusion of *Exhibitions, People* and *Annotations.* It uses controlled vocabularies, largely from the Getty Art and Architecture Thesaurus.

While our database design is built on standard relational structures, what makes it different to other designs is that its metadata is not explicitly built into tables. Instead a *meta-metadata scheme* is used to allow evolutionary change to metadata structures over time.

The interface is, and remains, changeable. This is significant if we consider that the social processes involved in supporting different forms of interface each pose requirements on a gallery's interface design. Further, these requirements are largely unknown to us. The literature on digital archives relating to experiences with the social processes supported by actual gallery systems is rare. One study by Bishop (1999) that compared the use of digital libraries by people from differing socio-economic groups found little effect. Users from diverse groups had the same tendency to be deterred by problems in the interface or inappropriate presentation of the library contents.

What we found is that social processes tend to dominate both system design and deployment. A shortfall of several gallery efforts of which we are aware (including our own prior works) is that they have experienced difficulties in the growth and management of collections over time. We have trialled several processes of Resource and Exhibition acquisition and have developed conjectures of social processes that may both benefit from and sustain gallery collections. The processes we have trialled (and continue to trial at the time of writing) include both small-scale *community-based* contributions and larger-scale *archiving* processes.

We used a gallery metaphor. Deliberately abstract, it was designed as a generic solution to support slide collections, lectures, student work-posting, online exhibitions of work and peer-review processes. For each of these processes we have conjectures of how the gallery might be deployed so as to sustain the development of significant collections of material over time.

As an interface, we chose the open-source implementation environment, *TikiWiki* (TIKI, 2004) As an open source initiative providing tools for building portals, it provides a relational database structure for supporting such common tools as articles, Blogs, FAQs, Forums and Wikis, interface building tools and user management. We have extended Tikiwiki with our relational structures. We have named our implementation with the non-acronym A•VI•RE (á-vēr´-ā).

4. USER SCENARIOS

The design of A•VI•RE is best explained through a description of its user scenarios. At the topmost level is the gallery system itself, which contains all objects (Exhibitions, Resources, Works, People and Annotations) related to the gallery. All users interact with A•VI•RE Exhibitions. At its simplest level, an Exhibition can be any user's collection of selected Resources. An Exhibition can be created by re-using existing gallery Resources, or by adding new Resources to the gallery, or a combination of both. Exhibitions comprise a subset of all the Exhibitions in the gallery, a subset of all gallery Resources and an interpretation (possibly null) of those Exhibitions and Resources.

The re-use of existing resources is a significant difference between A•VI•RE and many other digital libraries aimed at visual material.

4.1 Viewers

A Viewer is the generic gallery user and interacts with A•VI•RE at its simplest level. A viewer may browse or search the gallery contents. Although effective browsing and searching are by no means a trivial issue (Furnas and Rauch, 1998; Spink, et al, 1998), in the A•VI•RE model, this is the level of interaction that is open to all.

A Viewer enters the site by URL and may proceed to access the search function. The search function is offered in the usual manner. A form box, clearly labelled as the search tool, is presented. As with all such open Search functions, this level of interaction is extremely flexible yet effectively limited if the user has little idea of appropriate search terms.

The Browse option is the first indication to a novice user that A•VI•RE presents special functionality. It is here that a user is presented with the choice to browse Resources, Works or Exhibitions. Browsing Resources allows access to the entire raw, uninterpreted content of the library, with various starting search terms. These terms are presented as pull downs and are a list of categorised metadata.

Browsing *Works* is the next level of abstracting the library content. Thus, for a user to select (or enter) the name of a Work, such as 'Eiffel Tower' for example, they are asking the system to search and display all Resources

and Exhibitions that have Eiffel Tower as the name of the Work. This could include digital images that contain the Eiffel Tower, interpretations of the Eiffel Tower provided by gallery users, drawings of the tower, related landscape images or drawings, reference to books, paintings and sculptures, etc. and reference to related web sites. It could also return textual commentary (Annotations) if any of these resources have been the subject of interest to a critic.

The Exhibition is the primary abstraction within A•VI•RE. A Viewer browsing Exhibitions is presented with the choice of all Exhibitions, although the home page offers immediate access to the most recent or featured Exhibitions. An Exhibition is a subset of the gallery Resources, selected by an Exhibiter based on some point of interest.

A viewer could select from the list of Resources presented and title that selection 'My favourite views of Paris'. This Exhibition would not be public and would persist for the duration of the session. However, if the Viewer has a more abiding interest in the Resources, s/he can choose to register on the site (become a *Person*) and submit an Exhibition.

4.2 Exhibitors

As an Exhibitor, a user can not only access the gallery Resources, but also add their own level of meaning to those Resources, beyond that provided by the current metadata associated with each Resource.

For example, an academic in architecture could choose to select or add to the gallery a series of Resources that form the content of an upcoming lecture. If they choose to add Resources to the gallery, they are obliged to complete the metadata form associated with each Resource (see Fig.8). and to name this Exhibition. They would indicate on the submission page that this Exhibition is 'public', and 'not for peer review', adding an Annotation, which would contain a discussion of each of the Resources for the benefit of the students. This annotation could also refer them to other Exhibitions of interest, or suggest further research on the topic. By using A•VI•RE in this way, an academic has made a collection of visual Resources readily available to many groups of students.

This Exhibition persists. At some later time, the Exhibitor could choose to edit the contents of the Exhibition, perhaps updating some Resources or adding further comment for a subsequent lecture. The academic could create many such Exhibitions, in effect, a persistent gallery of work related to their teaching.

A•VI•RE also offers the opportunity to publish Resources in the traditional sense. By creating an Exhibition, adding Resources and metadata to A•VI•RE, adding a significant commentary, and submitting for 'peer review', one can access the eJournal function of A•VI•RE. This submission process makes the Exhibition available to an initially restricted audience. Online review, comment and revision are generally well understood. Hence, when (and if) the Exhibition is accepted, it is made publicly viewable. Once 'published', the Exhibition can be commented upon by any user. These additional Annotations are also available to all who wish to view the Exhibition. By this means, published works are open to ongoing commentary. This commentary could also include suggestions for adding to, or correcting, the metadata associated with any Resource.

As Resources can be part of an unlimited number of Exhibitions, the Resources that are submitted for a 'paper' can be subsequently used by others in related debate. That is, others may create other Exhibitions, using the same or similar Resources, in order to further interest in the topic.

Another category of Exhibitor is that of practitioner. As a creator of works, a practitioner may chose to document their work, and submit to A•VI•RE. In this and other forms of Exhibition, the curatorial process becomes important.

4.3 Curators

The curatorial process is at the heart of this gallery. In all cases, any A•VI•RE Exhibition must undergo some form of curatorial review.

All submitted Resources can only be made public after vetting. This vetting may be lightweight, for example, in the case of the academic submitting lecture material. For practitioners, a modest level of peer review may suffice, as the practitioner's clients are the effective reviewers of the work. For academic publication, a full peer review, blind referring process should be the norm.

A Curator may also be invited to select from the Resources of the gallery and present a themed Exhibition, much as an invited curator would do in a traditional gallery or archive. Such invitations could also include the submission of further Resources to A•VI•RE. In this way, significant collections of digital material can be curated and made public without the overhead of site creation or other technical considerations. In this light, A•VI•RE can be seen an international gallery, offering a large and geographically distributed community of interested participants the opportunity to engage.

4.4 Critics

The attached annotation function, which applies to every Resource and Exhibition, allows for easy and open critique. The Annotation function and its persistence within the gallery offer the opportunity for seamless access to previous comment on any Resource or Exhibition, providing potential historical interest.

Our implementation environment, TikiWiki, is an open source environment for building collaborative websites. It provides numerous common Internet tools that are said to aid collaboration, for instance, blogs, wikis, and forums. It is seductive to use these freely to support the Critic role in A•VI•RE. This we are resisting for methodological reasons. We want to understand the social roles that critics can play in online galleries. To do this, we are introducing functionality for Critics carefully. Annotations, an online analogue to the modern "sticky note" and the ancient marginal note, are our first, minimal move. Our next, and this remains in front of us at the time of writing, is to generalize Annotations to be *wikis.* The wiki is a free-form collaborative hypertext tool with the special features of induced page creation through *wiki names* and automatically maintained backlinks (PPR, 2004). There is considerable evidence to show that wikis are productive environments for small-group authoring around a set of issues, but little practical experience with them in our context of visual libraries.

5. THE METADATA SCHEME

All five major object types (Exhibitions, Resources, Works, People and Annotations) in A•VI•RE have metadata associated with them. This is a necessity for retrieval and search. We started the project with the assumption that we would use an established metadata scheme to store these data. We soon concluded that this would not be practical. The extant schemes we reviewed included Dublin Core, IATF-CDWA, VRA Core v3.0, Digilib and the scheme used in an online slide library at Melbourne University (the latter two were off-line at the time of writing).

The Dublin Core (DCMI, 2004) is a world-wide generic metadata standard. It focuses on capturing properties of online resources. As a global standard used as the basis of many other standards, it provides a *de facto* authority on resource-based data worthy of being captured and on preferred formats for such data. It does not, however, capture relations among resources, or relations among resources and objects in the world. The Dublin Core is used as a substrate on which to develop a new, relational standard within A•VI•RE.

Produced by the Art Information Task Force (AITF), the Categories for the Description of Works of Art (Getty 2004) are guidelines for formulating the content of art databases. They articulate an intellectual structure for descriptions of objects and images: in this sense they constitute a schematic representation of the requirements and assumptions implicit in the practice of the discipline of art history. They are presented specifically as guidelines expected to be adapted to the needs of specific collections. Their focus is on works of art *per se* – they provide ways to refer to information *extrinsic* to a work of art, but do not make these relationships sufficiently clear for the direct construction of a metadata scheme that requires them.

The Visual Resources Association VRA Core Categories V3.0 (VRA, 2004) presents a single scheme that can be applied to both works and images. In object-oriented terms, it applies polymorphism to field names. While initially we found this a compelling idea, its polymorphism was subsequently found to be dissonant. For example, the field SUBJECT is defined by the VRA Core as holding *terms or phrases that describe, identify, or interpret the Work or Image and what it depicts or expresses* (VRA 2004, subject). These terms are sufficiently distinct between Works and Images (analogous to our Resources) such that we found we should specify separate structures for both Works and Resources. Our Exhibitions, which are key to A•VI•RE, are not, however, supported by the VRA Core. That said, our resulting structure owes much to the VRA Core and can be seen as one that attempts to express as much in common amongst object types as is meaningful in our domain. Our use of common data can be most clearly seen in the A•VI•RE division between metadata *about the container* and metadata *about the container's contents.* Metadata *about the container*, shown in *Fig. 1*, which is common across all A•VI•RE types, records data describing the database record itself, such as a unique ID, the cataloguer,

the custodian and creation/modification dates. Metadata *about the container's contents* shares several fields across record types, with some key differences. Notably, we did not specify a TITLE field for resources. In all actual systems we reviewed, the practical use of the TITLE field for images/resources was to record a combination of terms from other fields, for example. WORK.title, WORK.creator and WORK.subject.aspect. Further, Resources and Works share many subfields of their subject field (with key differences in intended meaning), but Resources have the subfields *.taken, .aspect* and *.feature* that are distinct from those recorded for Works. We note that, throughout implementation, the metadata structure for Works, Resources and Exhibitions underwent significant refinement. It may well be that a review against the VRA Core V3.0 will, in time, reveal additional refinements making A•VI•RE and VRA Core more consistent.

The *Digilib* repository was an earlier project done by two of the authors (Docherty and Szeto) at the University of Queensland. We used its metadata scheme as a source of controlled vocabularies. We used the scheme from a now off-line repository at the University of Melbourne in a similar fashion.

Though not a metadata scheme in its own right, the Getty Art and Architecture Thesaurus (Getty, 2004a) is an authoritative compendium of terms relating to art and architecture, from which we sourced certain key terms. Though authoritative, it is not inclusive. For instance, none of *grammar, hierarchy, figural space, translational symmetry*, nor *scale symmetry* is in the AAT, yet all were relevant to us.

The lack of inclusivity of the world's foremost authority is but one signal of a larger issue. There is little question that metadata suffers a tension between service for search and retrieval across broad communities, and making meaningful distinctions within a community. Metadata will differ across communities and repository designs need to accommodate this reality. This is a well-known problem and is the root motivation for many current efforts on ontology-sharing in repositories.

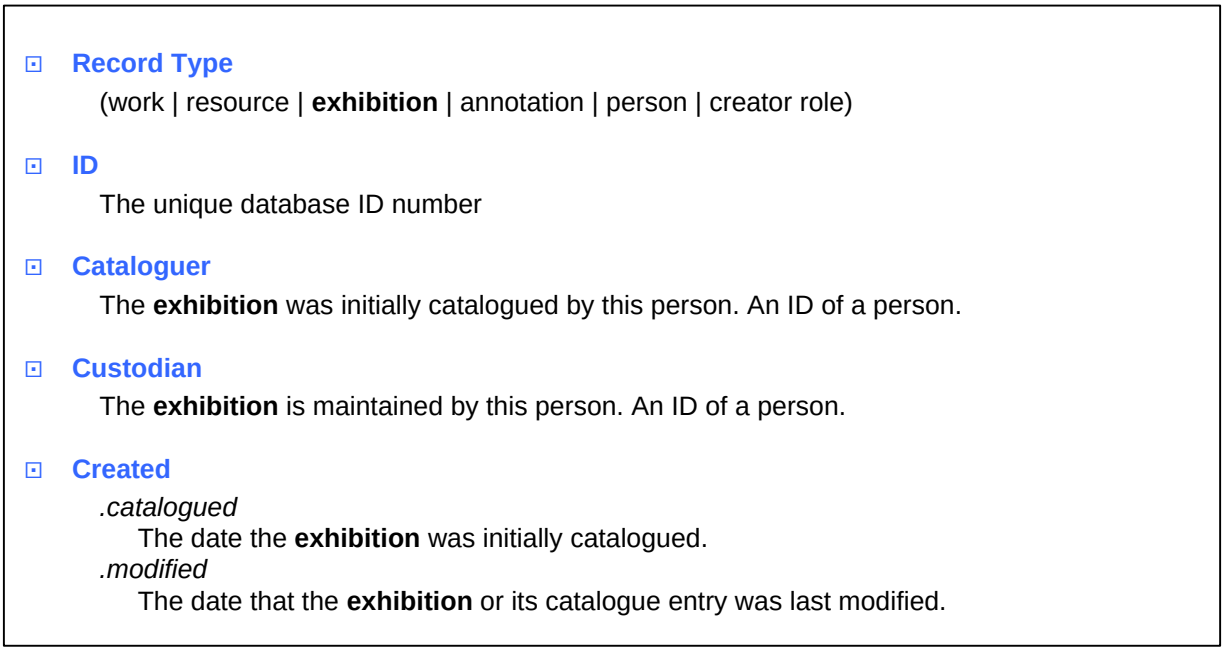

FIG. 1: A•VI•RE metadata about the container.

The problem for A•VI•RE was a precursor to ontology-sharing. We needed a way to modify metadata structures to some extent, without having to alter the database schema to do so. Our response lies in the design of a *meta-metadata scheme* in which the repository can accept different metadata schemes depending upon the collection being catalogued.

6. THE META-METADATA SCHEME

The process oriented description (Use Cases) of the A•VI•RE Repository and the metadata content specified by the A•VI•RE metadata scheme were used to inform the specification of the A•VI•RE database. The methodology used to specify a data-centric view is called Object Role Modelling (ORM) (Halpin, 2001). ORM is a structured high-level data modelling approach based on the expression of information as elementary English sentences

(elementary facts). These elementary facts are then explicated using a formal language, which can then be transformed to a logical model (tables and relationships) and finally into a physical design (database level). This methodology is a simple yet powerful way to design and build databases.

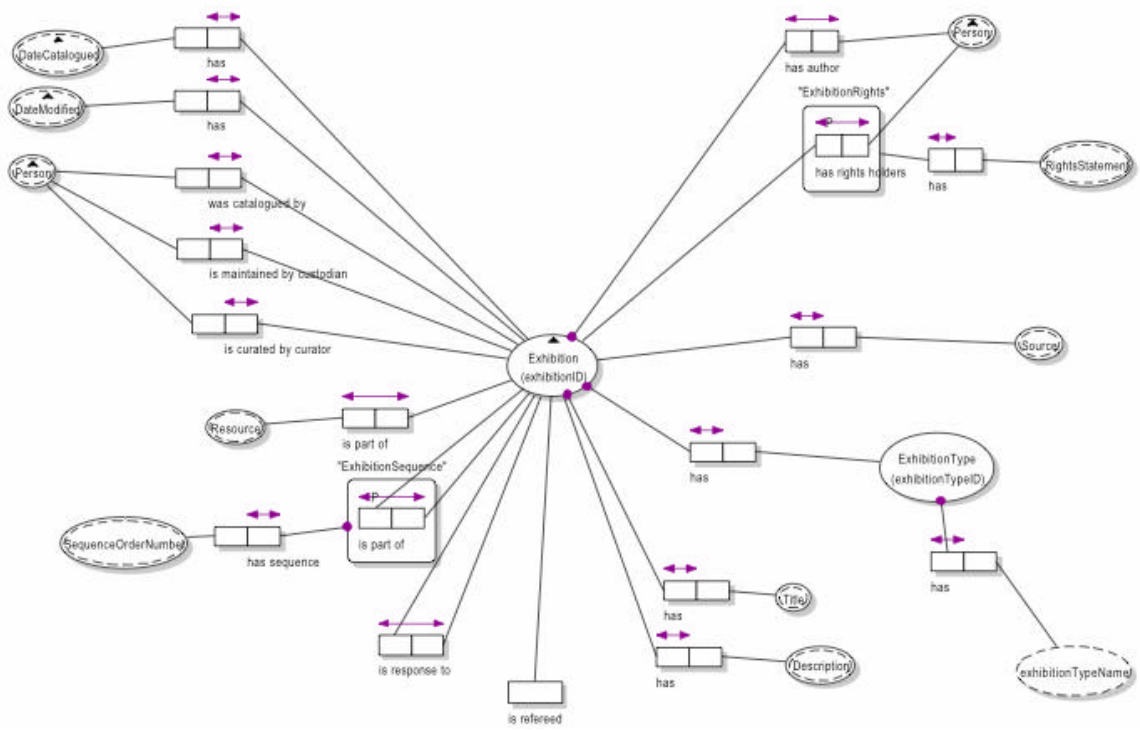

FIG. 2: Object Role Model diagram capturing facts about an Exhibition and its relationship with Resources: "Exhibitions may contain many Resources" and "Resources can be a part of many Exhibitions".

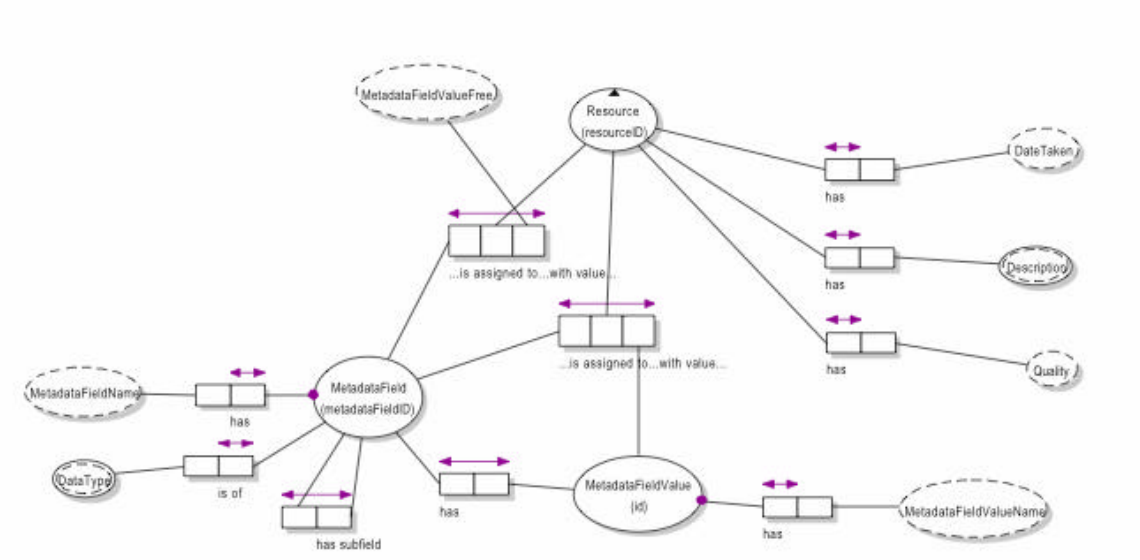

FIG. 3:. Object Role Model diagram modelling generic metadata (both controlled vocabulary and free-text) associated with Resources.

For example, in Fig. 2, the ORM diagram depicts the fact that Exhibitions have a Title, Description and ExhibitionType. Using ORM, many facts about A•VI•RE were modelled but are too numerous to be shown in full here.

The advantage of using a high-level conceptual modelling system is that it is possible to model the metadata scheme itself. The metadata fields (Aspect, Feature, Form, Structure, Material and Performance) can be explicitly stored in the database as "hard-coded" field names. The disadvantage with this approach is that the A•VI•RE metadata scheme is "hard-coded" into the database and changes to the scheme also requires modification of the database and the programming logic. Another approach is to model metadata generically. The ORM model in Fig. 3 depicts "MetadataFields" (e.g., Aspect) having a controlled vocabulary of "MetadataFieldValues". These fields and values are then combined with a Resource to store the fact that this Resource has metadata with a certain value. An example ORM elementary fact is: "The MetadataField 'Aspect' with ID 1 is associated with Resource with ID 1641 with a value of 'Perspective' for MetadataFieldValue." Similarly we can also have: "The MetadataField 'Features' with ID 6 is associated with Resource with ID 1641 with a value of 'Arched Windows' for MetadataFieldValue."

By storing the names of the metadata fields within a record in a table rather than as field names, future modifications to the schema can be supported. Another advantage with this generic approach to metadata storage is that other metadata schemes such as Dubin Core or IMS can be supported within A•VI•RE. We have loosely referred to this as the meta-metadata scheme.

7. IMPLEMENTATION

As the implementation environment, TikiWiki can be used to create Web-based collaborative applications with modules such as Wikis, Weblogs, forums, chat rooms, polls, surveys, and user and group management. Web sites, Intranets, extranets and web portals can be built also using TikiWiki's Web-based administration interface.

FIG. 4: Initial prototype of the A•VI•RE Web Application, using the basic implementation of TikiWiki.

Complete customisation is also possible because TikiWiki's PHP (PHP, 2004) source-code and relational database is made available. The application architecture makes a clear distinction between the data, programming logic and presentation logic with a templating engine and Cascading Style Sheets (CSS). This supports a separation of the development roles in Web development projects: database design, programming, and user experience. For example, a user-experience engineer can develop new themes while a software engineer is programming a new TikiWiki module.

The A•VI•RE application has been implemented using a subset of the TikiWiki modules and incorporating the A•VI•RE data model into its relational database. User collaboration features such as Wikis, Forums, chat rooms, Shout box, polls, RSS feeds and Weblogs form the community aspect, while the Exhibition functionality has been implemented by programming our own custom TikiWiki modules. TikiWiki also provides Image Gallery and File Gallery modules, which provide a subset of the functionality required by the Exhibition metaphor. For example, Images can only belong to one Image Gallery, while in A•VI•RE, Resources can belong to more than one Exhibition.

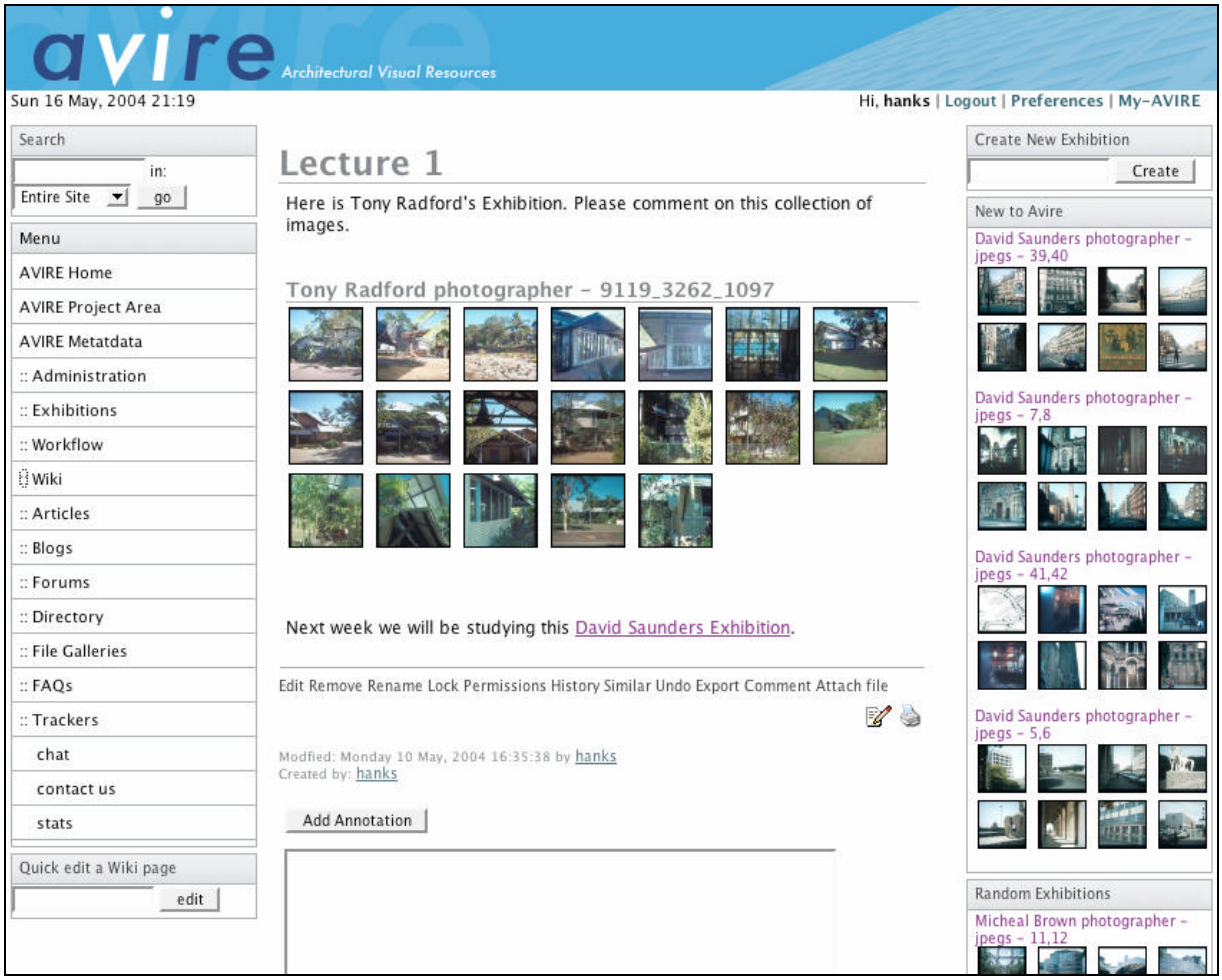

FIG. 5:. Release version of the A•VI•RE Web Application, based on TikiWiki. Note the tools for community building on the left, in particular the Forums and Blog tools.

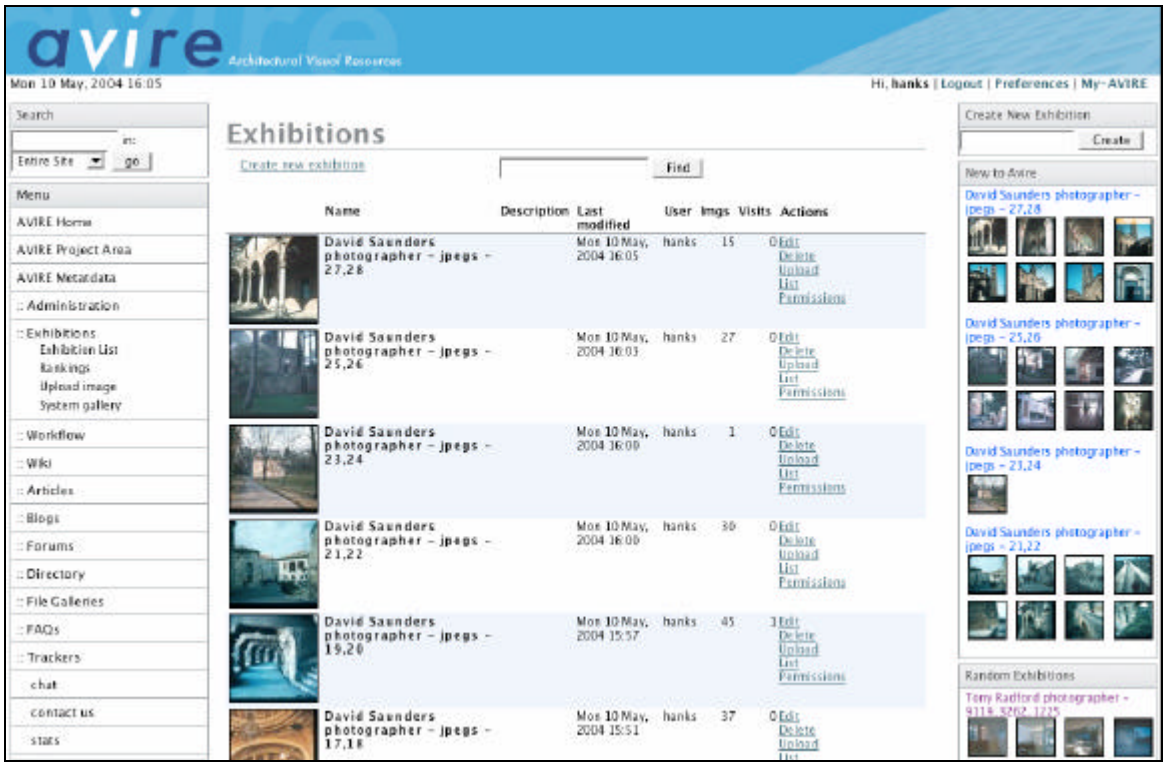

FIG. 6:. A list of Exhibitions.

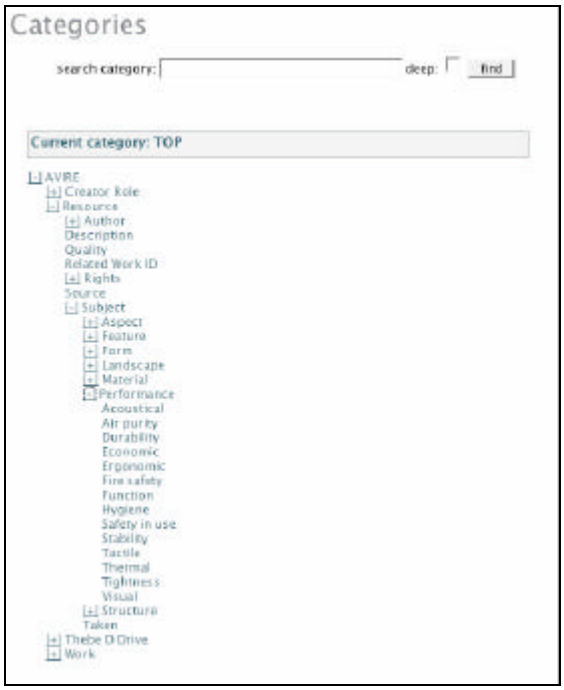

FIG. 7: Browsing Resources by metadata. Users can browse the AVIRE metadata hierarchy for Resources using an expandable tree interface.

8. META TAGGING AND AUTHORING EXHIBITIONS

A•VI•RE is intended to serve a broad range of user communities, from students submitting assignments to instructor-managed galleries, to supporting scholarly journals. A main (and yet unverified) hypothesis of A•VI•RE is that resource sharing across a range of academic and cultural discourse is an effective way to build vibrant and self-sustaining collections. A•VI•RE must thus appeal to many kinds of authors and many kinds of authoring acts. We have designed and are implementing four different authoring styles for Exhibitions. These represent distinct alternatives within the largely unexplored design space of Exhibition authoring. They are as follows:

1. "Slide Table" – web form upload and meta tagging of Resources.
2. Wiki – sequence interleaving of text, Resource, text.
3. Mass import including HTML – importing a web page.
4. Structured paper format – scholarly writing.

8.1 "Slide Table" authoring

"Slide Table" authoring is named after its analogy to the physical practice of organizing collections using a backlit slide table. It is a two-step process. In step one, a zipped archive of Resources is uploaded to the Web server, which is then unpacked and stored in the A•VI•RE database. The final step involves metatagging; a web form is presented, requiring the user to add metadata to the Resources before the Exhibition can be published. To speed this process, bulk tagging of Resources is supported, by allowing users to apply previously selected metadata and also applying metadata to multiple-selected Resources.

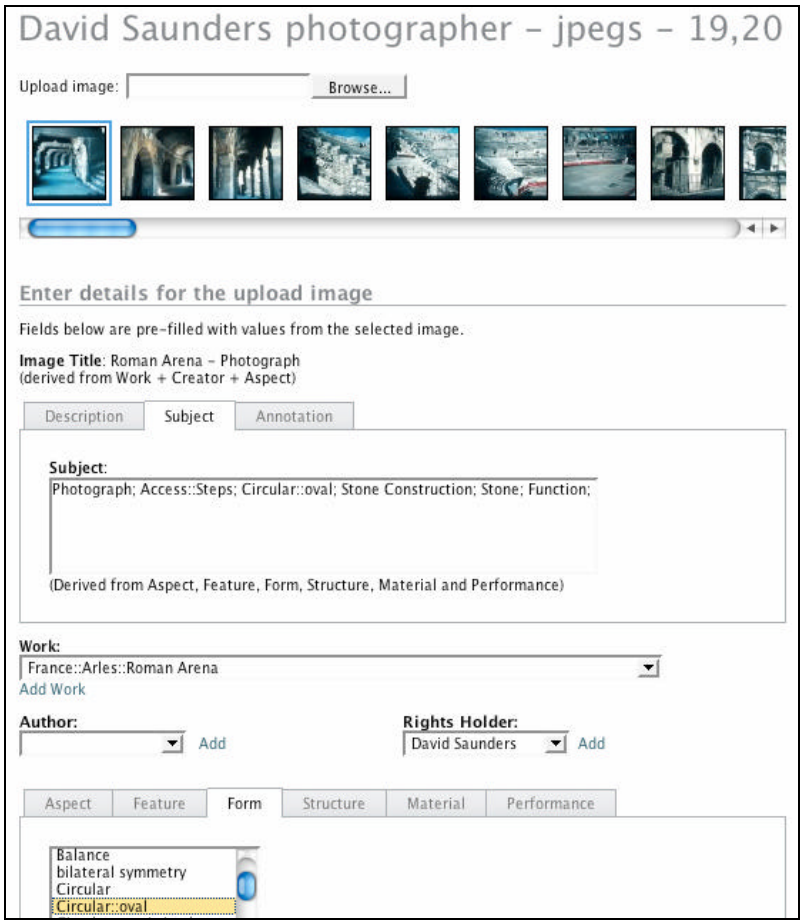

FIG. 8:. Slide Table authoring of Exhibitions. Resources are uploaded and then tagged. To ease data entry, fields are pre-filled with data from the previous uploaded Resource. The Image Title is derived from the Work Title, Creator and the Aspect. The Subject Field is automatically derived from the concatenation of Aspect, Feature, Form, Structure, Material and Performance metadata fields.

8.2 Wiki based authoring

The original Wiki was developed to support idea sharing and documentation about best practices or "patterns" in software engineering through web-based authoring (PPR, 2004). Anyone is allowed to edit any web page, relying on content versioning and the community of users/authors to ensure that contributed content added value. Since 1995 as many as 300 different Wiki clones have been developed in various programming/scripting languages with various levels of functionality (WIKIENG, 2004). The most popular extension in many of these implementations is applying a permissions layer over the Wiki functionality so that only authenticated users are able to edit pages. In some implementations such as TikiWiki, Wiki functionality is a module that sits alongside other modules such as forums, Blogs, Polls, News and image galleries.

The Wiki suggests an interesting way of authoring small exhibitions, conceived of as text with interleaved images and other resources. The present Wiki almost achieves the user-interface aspect of this task. A Wiki author can insert a file into a Wiki, but the insertion must be done at the moment of writing and requires one publish cycle to do. It is better to be able to defer the insertion of objects. This can be done by having two reserved WikiWords: "AttachFile" and "AttachResource". Inclusion of an "AttachFile" in a WikiPage results on publication of the page in a button. Clicking on the button opens a local file browser. Inclusion of an "AttachResource" in a WikiPage results on publication of the page in a button. Clicking on this button opens a browser into A•VI•RE.

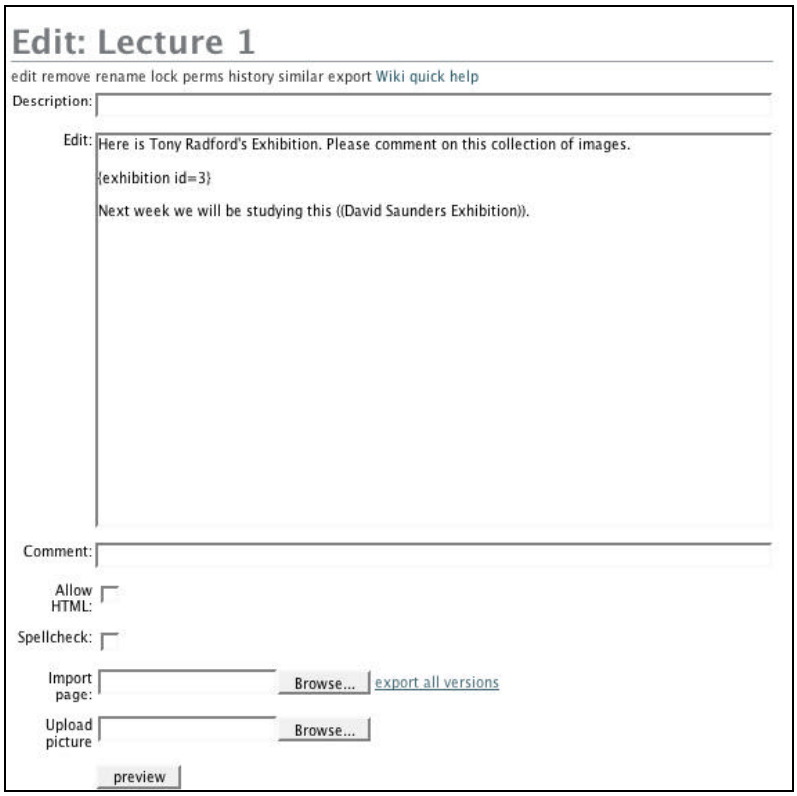

FIG. 9: Exhibitions can be placed within Wiki pages by referencing the Exhibition's identification number. Wikis were designed to simplify the process of creating web pages without requiring in-depth knowledge about HTML. The displayed exhibition can be seen in Fig. 5.

Another extension to the Wiki authoring environment is the ability to reference and dynamically include Exhibitions and Resources, which are displayed in-line within the Wiki text. For example "{exhibition id=3}" would retrieve Resource thumbnails from the Exhibition with an ID of 3 and display it amongst the Wiki text (Fig. 5 and Fig. 9). Similarly, "{exhibition latest max=3}"would display the latest three Exhibitions (Fig. 12). Some other examples of extensions to the Wiki authoring language are:

- {resource id=1641} – displays the resource with id of 1641
- {exhibition title="lecture1"} – displays the exhibition with title "lecture 1"

- {exhibition author="David Saunders" latest max=3} – displays the last 3 exhibitions by David Saunders.
- {resource work="Sydney Opera House" latest max=10 quality=5+} – displays the latest ten resources on the Sydney Opera House that have a rated quality better than 5.

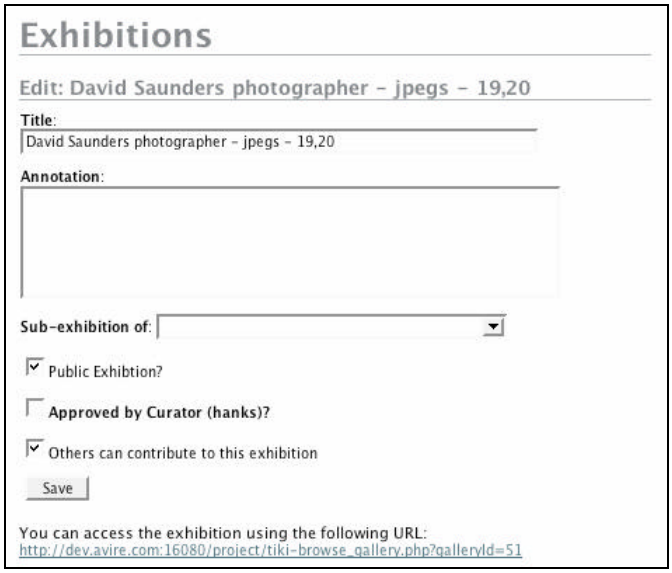

FIG. 10: Edit interface for Exhibitions. Exhibitions undergo a curatorial process. Once a curator approves an Exhibition, it is visible to the public. An exhibition can be a sub-exhibition of another exhibition.

FIG. 11a:. Browsing Resources in an Exhibition.

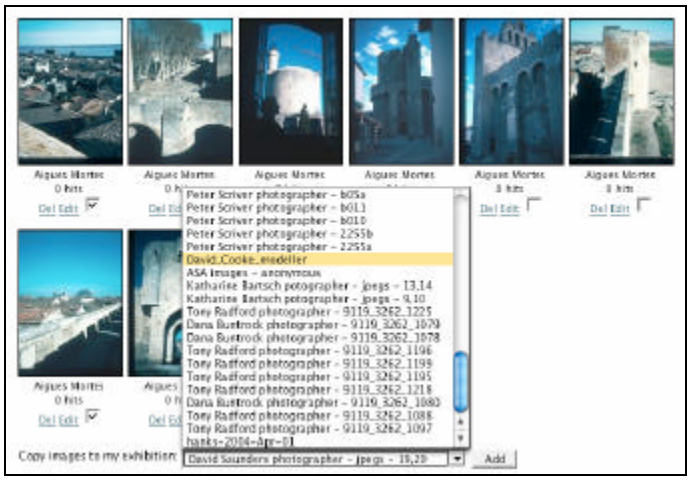

FIG. 11b: Exhibitors can also create Exhibitions using Resources created by other Authors. This is achieved by selecting the desired Resources and adding them to the user's own Exhibitions.

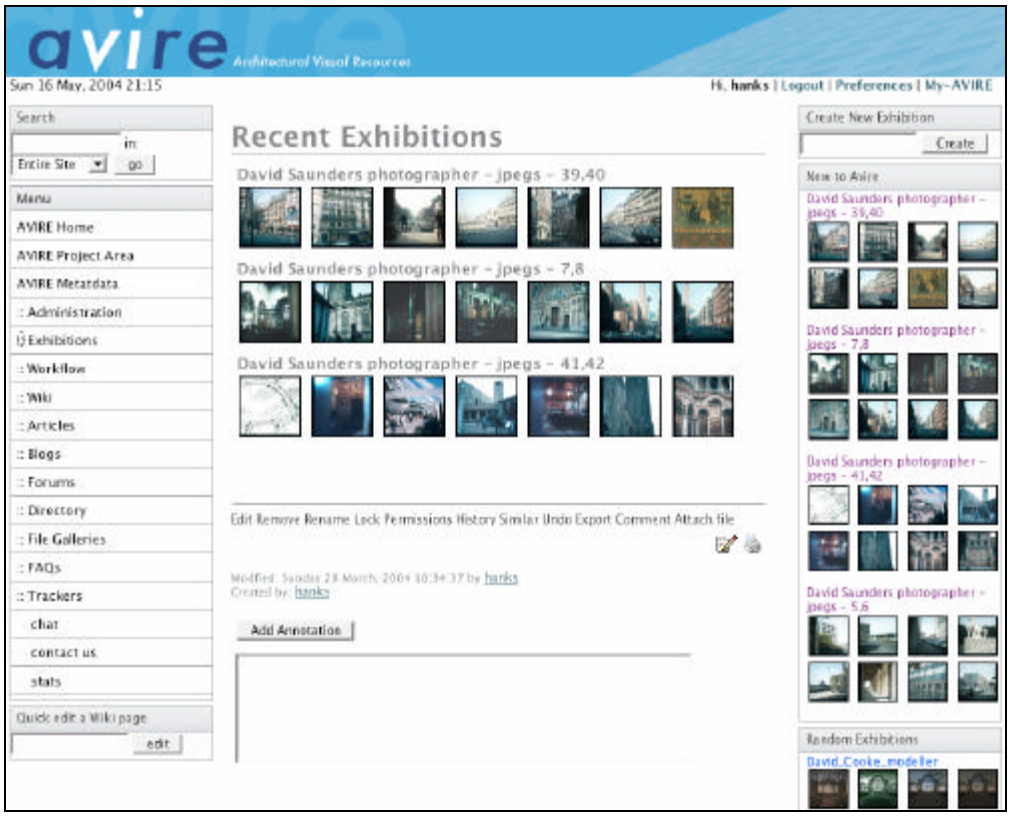

FIG. 12: The most recent three exhibitions are retrieved and displayed according to the following enhanced Wiki authoring language: {exhibition latest max=3}

8.3 Mass import

Import of collections into A•VI•RE is through a text file that specifies a local path to each object and the metadata associated with the object. As a convention, collections are stored locally within a single directory (and its subdirectories). The text file (which will have a specified name) is stored in the same directory. The import process consists of creating the import text file, transferring its containing directory (and contents) to the A•VI•RE server and running an A•VI•RE import script.

While the text file is human-readable and editable, it is primarily meant as a target file for a cataloguing system. Most online repositories implement online metatagging systems. Online metatagging is a significant bottleneck in the process of developing online content. Desktop cataloguing systems are often more sophisticated, supporting functions such as drag-and-drop interactions, and more importantly, with the ability to deal simultaneously with sets of resources. There are several commercial asset management systems available on the market, all of which have some form of cataloguing facility. The system we chose is CUMULUS V5.5 (CANTO, 2004), which runs on a variety of platforms.

Fig. 13 shows the CUMULUS interface. A cataloguer works in CUMULUS by placing file icons corresponding to resources into categories represented as folders. Resources can belong to more than one category. The structure of categories is represented in a human-readable and editable text file. Mass import provides a way to catalogue entire websites and thus a familiar authoring environment for exhibitions. Exhibitors can create websites interpreting a collection of resources. The website is then catalogued, with its individual images becoming Resources and its main files becoming Exhibitions. The minor elements of a website, such as graphic buttons, may be left without metadata. In this case, they become Resources, but are findable only as they appear as an explicit part of an Exhibition.

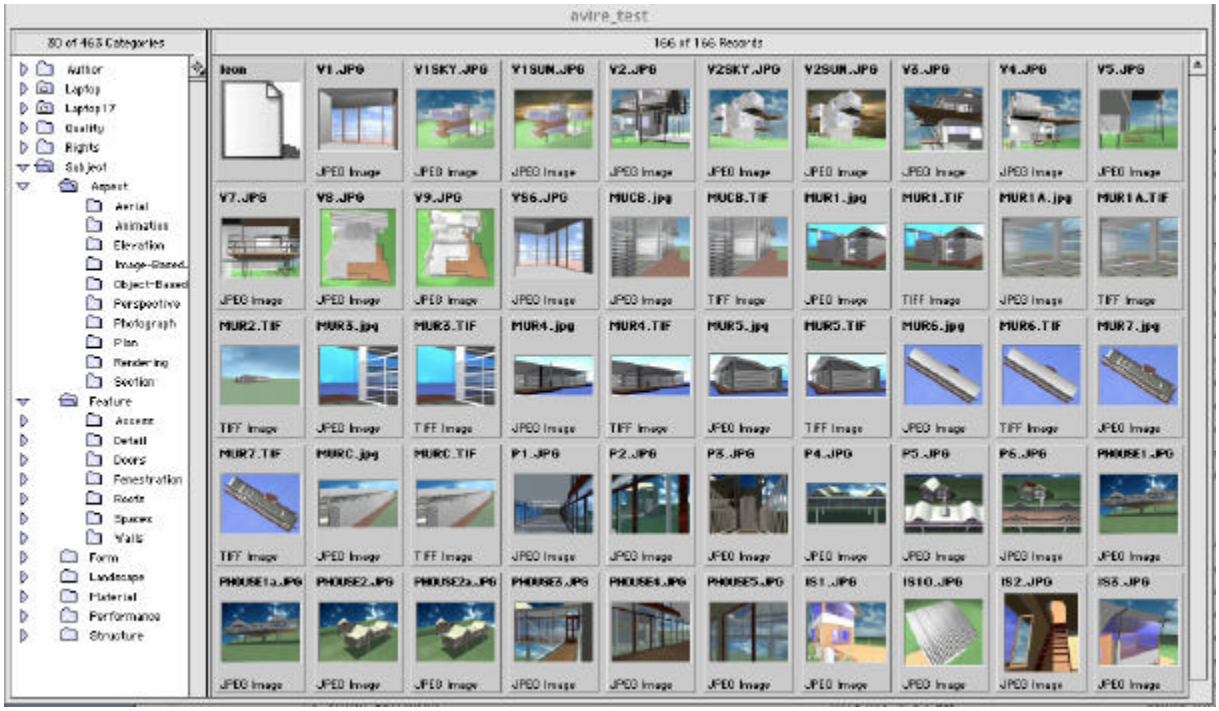

FIG. 13: The CUMULUS interface. CUMULUS generates and displays thumbnails of resources (shown on the right). A cataloguer can read in a set of categories, each representing a value of a metadata field, from a text file. These appear as a set of folders (shown on the left). The cataloguer can drag and drop sets of thumbnails onto categories or sets of categories onto thumbnails. The result is that the resources become associated with the categories. Thumbnails can be associated with several categories.

8.4 Structured paper

Scholarly writing is typically done against a paper format. In GUI-oriented environments (such as Word or Framemaker) the format is typically expressed as a template. In LATEX, it comes as a style file: part template, part macro language. In XML, it comes as a DTD or XML Schema that expresses conventions of scholarly writing in its tags. We intend to use the GCAPaper (GCAPAPER 2003), a widely used conference paper format. Resources within a GCAPaper formatted Exhibition would be metatagged using either the Slide Table or Mass Import processes. At the time of writing we have just begun the process of supporting GCAPaper-structured Exhibitions.

9. RESULTS AND SUMMARY

Our gallery system exists as an implemented, online system. It currently has significant private foundation funding in Australia and government funding in Canada. In early 2004 we plan a first public release. At that time the gallery will contain at least 3000 high-quality images of architectural works from Australia, Canada and other locations in the world.

We plan to support A•VI•RE from at least two sites, one in Australia (RMIT University) and one in Canada (Simon Fraser University). It will be available from http://www.avire.org.

A•VI•RE presents a design for online galleries and a hypothesis, both practical and research. The hypothesis has several components as follows:

- that a gallery metaphor structured around Exhibitions, Resources, Persons, Works and Annotations and accessed by the roles Curator, Exhibitor, Critic and Viewer can be used to support a range of authoring process;
- that Resource and Exhibition sharing across Exhibitions within A•VI•RE can provide significant reuse of resources;
- that the value, to authors and their audiences, of being able to put interpretations of collections into the system as primary objects can yield significant motivating value for using a gallery system.

We based these hypotheses on prior systems, notably vGallery (Shannon et al., 2001, Woodbury et al., 2000), in which community-based authoring around exhibitions was shown to be an enabling tool for several forms of online learning. Otherwise these hypotheses have not been tested. In front of us are two directions of research. The first is a process of design space exploration in which the A•VI•RE interface and the tasks for which it is suited are jointly developed. The second must occur within each significant design state of the first. It comprises evaluation of the A•VI•RE system against the patterns of use it will experience.

10. ACKNOWLEDGEMENTS

The research and development reported here was supported by the Ian Potter Foundation, the Heritage Canada New Media Research Networks Fund, the British Columbia Advanced Systems Institute, The University of Adelaide, The University of Queensland and Simon Fraser University.

11. REFERENCES

Anonymous, CANTO (2004). Canto website. accessed at http://www.canto.com on 4 Feb 2004.

Anonymous. DCMI (2004). Dublin Core Metadata Initiative. accessed at http://dublincore.org/ on 4 Feb 2004.

Anonymous, PHP (2004). PHP website, accessed at http://www.php.neton 16 May 2004.

Anonymous, PPR. (2004). Portland Pattern Repository's Wiki. accessed at http://c2.com/cgi/wiki?WelcomeVisitors on 16 May 2004.

Anonymous, UQ (2004). Architecture Department website, The University of Queensland, accessed at http://www.architect.uq.edu.au http://www.architect.uq.edu.au on 16 May 2004.

Anonymous, VRA (2004). VRA Core Categories. Version 3.0 a project of the Visual Resources Association Data Standards Committee, accessed at http://www.vraweb.org/vracore3.htm on 16 May 2004.

Anonymous, VRA (2004a). VRA Core Categories -- Subject, Version 3.0 a project of the Visual Resources Association Data Standards Committee, accessed at http://www.vraweb.org/vracore3.htm#subject on 16 May 2004.

Anonymous, TIKI (2004). TikiWiki Website.accessed at http://www.tikiwiki.org/ on 4 Feb 2004.

Anonymous, GCAPAPER (2004). Tools for XML Conference Presentations. SchemaSoft website, accessed at http://www.schemasoft.com/gcatools/ on 10 Feb 2004.

Anonymous, WIKIENG (2004). WikiEngines. accessed at http://c2.com/cgi/wiki?WikiEngines accessed on 4 Feb 2004.

Bishop A.P. (1999). Making Digital Libraries Go: Comparing Use Across Genres. Proceedings of ACM DL ’99, pp 94-103.

Furnas G.W. and Rauch S.J. (1998). Considerations for Information Environments and the NaviQue Workspace. Proceedings of ACM DL ’98, pp 79-88.

J.P. Getty Trust, Categories for the Description of Works of Art. (2004). accessed at http://www.getty.edu/research/conducting_research/standards/cdwa/index.html on 4 Feb 2004.

J.P. Getty Trust, Art and Architecture Thesaurus. (2004a). accessed at http://www.getty.edu/research/conducting_research/vocabularies/aat/ on 4 Feb 2004.

Halpin T. (2001). Information Modeling and Relational Databases: From Conceptual Analysis to Logical Design, Morgan Kaufmann, San Francisco USA.

Pitt S.P., Updike C.B. and Guthrie M.E. (2002). Integrating Digital Images into the Art and Art History Curriculum. EDUCAUSE Quarterly, vol 25, pp 38-44.

Shannon S. J., Roberts I. W. and. Woodbury R. F. (2001). vGallery: Scaffolding Reflection-in-Action for Students and Teachers, Proceedings of the 18th Conference of the Australasian Society for Computers in Learning in Tertiary Education, ASCILITE 2001, Dec 9-12, 2001, pp. 141-144.

Sharon P. Pitt, Christina B. Updike, and Miriam E. Guthrie. Integrating Digital Images into the Art and Art History Curriculum, EDUCAUSE Quarterly, 25 (2002): 38-44.

Spink A., et al (1998). Searching heterogeneous collections on the Web: Behaviour of Excite users. Information Research, Vol 4 No 2.

Witten I.H., Bainbridge, D. and Boddie S.J. (2001). Greenstone: Open-source DL software, Communications of the ACM, vol 5, p47.

Woodbury R., Roberts I. and Shannon S. (2000). vGallery: Web Spaces for Collaboration and Assessment, Australian University Teaching Committee, 2000 National Teaching Forum, Assessment in Higher Education, Canberra 6-7 Dec 2000, on the WWW http://www.autc.gov.au/forum/2000/papers.htm (as of 10 February 2004).