

- [Register](#)
- [Key Dates](#)
- [Schedule](#)
- [Local Info](#)

March 22-25, 2006
Albuquerque, New Mexico

Steve.museum: An Ongoing Experiment in Social Tagging, Folksonomy, and Museums

Susan Chun, The Metropolitan Museum of Art, Rich Cherry, Guggenheim Museum, Doug Hiwiler, The Cleveland Museum of Art, USA, Jennifer Trant, Archives & Museum Informatics, Canada, and Bruce Wyman, Denver Art Museum, USA

<http://www.steve.museum>

Abstract

Social tagging applications such as flickr and del.icio.us have become extremely popular. Their socially-focussed data collection strategies seem to have potential for museums struggling to make their collections more accessible and to build communities of interest around their holdings. But little is known about the terminology that visitors to museum sites might contribute or how best to obtain both useful terms and on-going social involvement in tagging museum collections. In the steve.museum project, a number of art museums are collaboratively researching this opportunity. These research questions and an architecture for a prototype research application are presented here. Prototypes created to date are discussed and plans for future development and term-collection prototype deployment are presented. We discuss the potential use of folksonomy within museums and the requirements for post-processing of terms that have been gathered, both to test their utility and to deploy them in useful ways.

Keywords: Social tagging, folksonomy, vocabulary, terminology, museum cataloguing, retrieval, open source, art documentation, social computing, accessibility

A. Why Social Tagging?

One year ago, the phenomenon of social tagging was relatively little known – limited to a small number of applications within popular, but narrowly-focused, Web sites such as flickr.com and del.icio.us. During the last twelve months, however, commercial Web sites such as Amazon.com and Google's gmail.com have incorporated elements of tagging into their user experiences, and user-supplied keyword indexing to improve access to a broad range of networked information resources appears to be well on its way to acceptance by both information scientists and users. Diverse communities – including music lovers (<http://www.last.fm>), Madonna fans (<http://www.madonna.com/taggingproject.html>), and gamers (<http://millionsofgames.com/>) – are forming and interacting via relatively simple applications that put the power of description and access into the hands of users. Museums have an opportunity to enhance user experience and to experiment with new strategies for user engagement that builds community and attract new audiences.

1. Improving Access to Museum Collections On-line

Museums collect objects and works of art that embody our cultural and natural heritage. Presentations about these objects (in print, on site, and on-line) recognize that contextual information about collections is often as important as the objects themselves: many objects “receive their significance only through the thoughts that cluster around them” (Boas, 1907). Museums have always communicated such information through exhibitions, publications, lectures, tours, and school programs. The emergence of the Web offered museums an opportunity to present more of their collections and context to a larger, more geographically-dispersed audience than ever before. The early years of publishing for this broader audience did not, however, result in a significant change in

the nature of the museum's engagement with its viewer: in the early days of the Web, museums continued to publish "curated," largely linear on-line exhibitions and educational materials.

Soon, however, museums began to take advantage of the on-line environment to expose their collections databases to Web visitors. Museums offered images and information about works in their collections – sometimes every work in their collections – through interfaces that were not much different from the database interfaces used by their own professional staffs. This model for presenting collections information, though it initially met with acclaim, has also proved only partially satisfactory. Works retrieved from collection databases lack context. Searching a database of thousands of works can be a frustrating experience for users unfamiliar with the idioms of art historical documentation.

We have come to realize that both models for presentation of on-line information about our collections – highly authored, linear exhibitions and educational materials, and un-interpreted collections databases – are inherently limiting. We've recognized Peter Walsh's "Unassailable Voice" in our writing (Walsh, 1997), whether it is the curatorial point of view in an exhibition or the pedagogical perspective of a lesson plan. And even "free-choice" interactive experiences are developed with an underlying and specific message in mind. At the other end of the spectrum, our collections databases isolate objects; scientific collections documentation is recorded in the specialist discourse of curators and registrars, a language unfamiliar to members of the general public, whose searches against that information often fail to yield results users expect or understand.

Simply put, the access offered by the Web hasn't translated into accessibility. This is one of the reasons that social tagging, where access points are supplied and shared on-line by the general public, may turn out to be an attractive solution to some of the problems of access. The folksonomy that results from social tagging appears likely to fill gaps in museum documentation practices that took their fundamental requirements from the business of museums (for example, Canadian Heritage Information Network (CHIN), [2005]; McKenna & Patsatzi, 2005). Tagging lets us temper our authored voice and create an additional means of access to art in the public's voice. For museums, including these alternative perspectives signals an important shift to a greater awareness of our place in a diverse community, and the assertion of a goal to promote social engagement with our audiences.

2. Proof Of Concept

The concept that cataloguers from the general public could augment museum professional cataloguing with valuable terminology in their own vernacular was tested with a small group of volunteers in an informal experiment conducted by The Metropolitan Museum in Fall, 2005. The volunteers were asked to supply keywords for 30 images from the Museum's collections. The terms supplied by the volunteers were compared to curatorial cataloguing from the Museum's collections management system, and the "unique" terms – new keywords not previously available through mining museum data – were "validated" for relevance to the work of art by a group of Museum staff members. For the 30 images tested, approximately 80% of terms submitted by the community cataloguers were unique – new additions to the museum's documentation. Example terms for Diana and the Stag (Figure 1) are shown in Table 1.



Figure 1: Joachim Friess (German, ca. 1579-1620). Diana and the Stag. First quarter of the 17th century (about 1620). The Metropolitan Museum of Art. Gift of J. Pierpont Morgan, 1917 (17.190.746).

Cataloguing from The Metropolitan Museum's Collections Management System

Sample Terms Collected

Artist/Maker(s)	Maker: Joachim Friess ca. 1579-1620, m. 1610	
Title	Diana and the Stag	A sampling of unique terms (not in the museum's collections management system) collected from volunteer cataloguers.
Object Name	AUTOMATON	
Date	First quarter 17th century (about 1620)	<ul style="list-style-type: none"> • action • animal • antler • archery • arrow • baroque • bejeweled • bow • bridle • chain • clockwork • container • deer/reindeer • dog/dogs • figurative • gilded • gold
Culture	German (Augsburg)	
Made in	Country: Germany City: Augsburg	
Medium	Silver, partly gilt, jewels, enamel	
Classification	Metalwork-Silver	

Dimensions	14 3/4 x 9 1/2 in. (37.5 x 24.1 cm)	<ul style="list-style-type: none"> • horse • hunt/hunter/hunting/huntress • intricate • jumping • leap/leaping • luxury • mechanical • movement • mythological/mythology • nude • ornamental/ornamentation /ornamented • ornate • precious • rearing • rider/riding • saddle • woman
Credit Line	Gift of J. Pierpont Morgan, 1917	
Department	European Sculpture and Decorative Arts	

Markings

[1] A pine cone (Augsburg town mark in use between 1620 and 1625). See Helmut Seling, *Die Kunst der Augsburger Goldschmiede 1529-1868*, Munich, 1980, vol. III, p. 19, no. 44. [2] Roman capitals I F within a rectangle (cameo) (maker's mark of Joachim Friess or Fries, born in Lubeck about 1579 master goldsmith in Augsburg in 1610, died 1620). See Seling, Vol. III, p. 143, no. 1248. [3] A Roman capital B below a crown and within a shield (cameo), unidentified. [4] A Roman capital E below a crown, and within a shaped reserve (cameo) (French guarantee mark for Foreign silver in use between 1809 and 1819). See Tardy, *International Hallmarks on Silver*, Paris, 1981, p. 190. Location of marks: [1] On the neck and inside the neck, and on the base. [2] Twice on the neck, on top of the base. [3] Inside and on the outside of the neck. [4] Inside the neck.

Description

Diana seated on stag, with Cupid on croup, and three hounds below; stag's head removable, body a flask; spring-driven mechanism in base permits automaton to move. Diana wears a crescent set with eight rose diamonds and one pink tourmaline; stag with coronet, pierced collar, lion mask on breast and trappings of scrolling ornament and cherubs; set with three pink imitation tourmalines; on base three lizards, two of them with green enamel, one small lizard, two frogs and four beetles, one of them with green enamel; octagonal casket-like base, decorated with applied ornament in same style as trappings; in center of left side movable key-hole cover on spring.

Accession Number 17.190.746

Table 1: Collections management system documentation compared to user-supplied terms

Though the data set examined in this and other such early tests at the Met is extremely limited, the terms collected are valuable enough that the Met intends to store and deploy them almost immediately. In addition, anecdotal evidence obtained through interviews with participants in the test indicated that volunteer cataloguers found the cataloguing task inherently enjoyable, despite the fact that the testing interface was not specifically designed to entice.

3. Tagging and Museum Communities

Museums want their communities to connect with their collections (Durbin, 2004). Our goal – whatever the method adopted – is to facilitate a unique and compelling engagement with a work of art.

Tagging appeals because it represents a dialogue between the viewer and the work, and between the viewer and the museum. A tag is a user's assertion that a work of art is about something. In the museum context, tagging offers a way for people to connect directly with works of art, to own them by labelling or naming them – one of the aspects of sensemaking (Golder & Huberman, 2005). Tagging also lets users assert personal connections. Recording these for future use makes re-discovery easier. Users remember the personal connection, rather than trying to re-imagine how the object might be discovered through a traditional search.




Figure 2: Tagging: Users assign tags to objects

For museums, the shared results of tagging offer chances to extend the experience in new and unexpected ways. It's easy to conceive of tag-powered visualizations and clouds (like Chudnov's visualizations of data from analog using Starlight (Chudnov, 2005)), "flythrough" navigations (like that of the Digital Depot in the Museum Boijmans van Beuningen, Rotterdam), or simply things that are just fun like flickr Tag Fight (Windisch, 2005 -). It also becomes possible to generate subscription RSS feeds of works of art based on a particular tag that can be easily shared. All of these ideas go towards creating a personalized partnership between the museum and the visitor.

Tagging can also become part of the museum's tool-set for fostering and maintaining interactions with teachers and students, or volunteers and docents. Here, rather than being motivated by personal gain (Vander Wal, 2005), social altruism, and other motivations for volunteerism in the cultural sector, play a role. This is reflected in the Cleveland Museum of Art's link to its on-line tagging tool: "Help others find this object" (Cleveland Museum of Art, 2005). It has also proven true at the Powerhouse Museum, where the Electronic Swatchbook project (Powerhouse Museum & Chan, 2005) is collecting terms in an application that supports indexing by volunteers.

What distinguishes tagging as a form of visitor engagement from other kinds of "interactive" museum programs is that the impetus lies not with the institution but with the individual; the visitor completes the experience. Tagging represents a personal investment in the museum's collection. Visitors add value for the museum, for themselves, and for other visitors by revealing distinct perspectives and connections, and recording them with tags.

B. Steve.museum

1. History

The participants in the Steve project came together following the Museums and the Web 2005 professional forum, "Cataloguing by Crowd" (Chun & Jenkins, 2005). The forum contemplated the possibilities of social tagging of museum objects and proposed the development of open-source tools to support term collection. The session was well attended by an engaged audience, and it was clear that there was interest in additional work on the idea. During the following summer, a week-long retreat – organized by The Metropolitan Museum of Art and the Guggenheim Museum and facilitated by Archives & Museum Informatics – was held on Grindstone Island, Ontario, to flesh out core concepts and define a project plan. The working group participants defined a number of potential uses for tagging at different museums, defined a shared research agenda (Cataloguing by Crowd Working Group & Trant, 2005), and created an initial specification for a front-end (interface) and a back-end (database).

Steve.museum is a collaboration to explore a number of issues associated with social tagging and folksonomic description. We're exploring the tools and techniques that support social tagging (and facilitate engagement with collections) and studying the resulting folksonomic terminology and its effectiveness in supporting improved access to museum collections. We're also interested in understanding what motivates individuals to contribute tags and in studying how this new sort of engagement with museum objects might help build new audiences and draw contributors who bring a multi-cultural perspective to looking at our works of art.

Steve is an open collaboration with an experimental methodology. Participants in steve come from many backgrounds, and work inside museums and in organizations that support them technically and intellectually. We pool raw data, resources, and research results in a distributed fashion that allows participants to move ahead on aspects of the project that coincide with immediate institutional needs, requirements, and abilities. The distributed execution of the project reflects diverse (and sometimes even conflicting) rationales for participation. We've realized that while some institutions are more interested in the social elements of steve, and others in the folksonomy, the same tools and methods enable the exploration of both.

2. Shared Research Agenda

We've identified a series of questions about social tagging and folksonomy in the museum. These questions are focused on getting, using, and understanding tagging data, and the potential social impacts of tagging. Central to this shared agenda has been an underlying research method that breaks down the different components of the experience, such as methods for building data sets, customizable front-ends, and plug-in architectures for the back-end and analysis. All of these different facets become variables that we can control and combine in different iterations of a tagging tool to see what experience best benefits from different permutations. This reflects our working assumption that while we may have different goals at the end of the day, the paths to reach those points require a common set of tools). It's within this shared research environment that we're exploring how best to use social tagging and folksonomic strategies for museum ends.

a. Systems Architecture Questions

Social tagging applications have evolved in two distinct ways. Tag servers, such as del.icio.us, citeUlike.com, PennTags, or dogear (Millen, Feinberg, & Kerr, 2005) store data separately from the source which is being tagged. A tag server can interface with a museum's existing data servers, and offer tagging on museum Web sites where on-line visitors already exist. Centralized systems such as flickr store tags and data on the same system, and require users to come there to participate. In our context this model would require either a) moving data from multiple museums into a single steve application, or b) deploying a totally distributed set of applications, comprised of local implementations of commonly designed tools.

We need to understand the pros and cons of each of these implementation paradigms and appreciate how they influence the experience for individuals, communities and the museum. Different implementation models are being explored in the prototypes now under development, with both a shared steve application and a single institution steve application now in beta (<http://www.steve.museum/>).

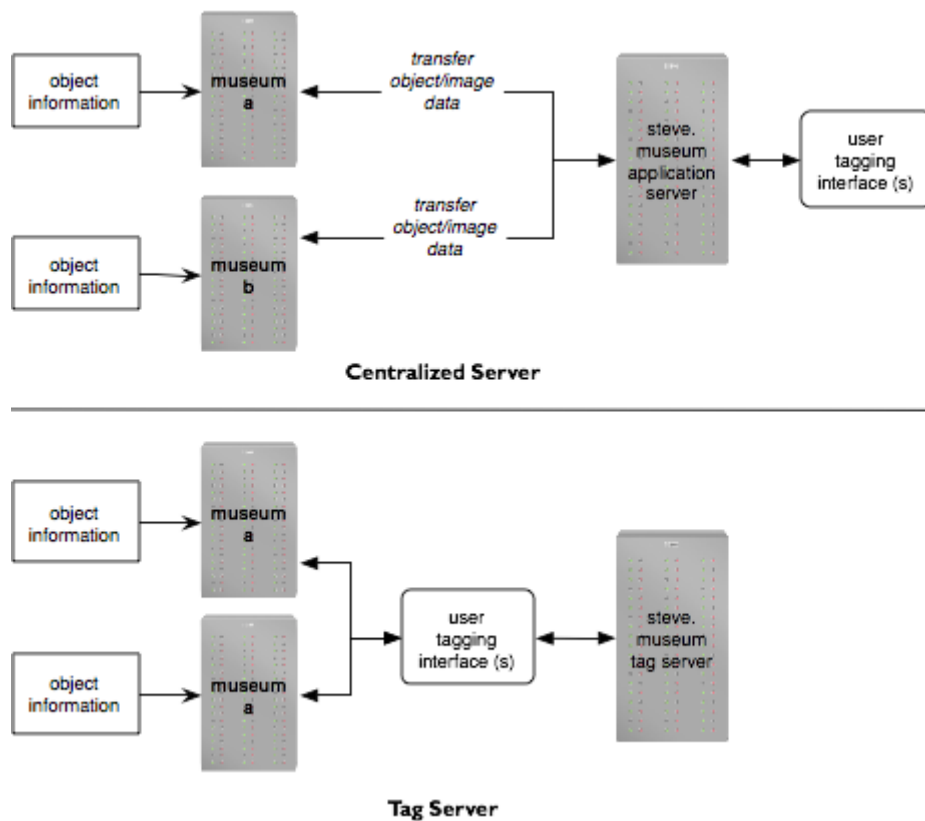


Figure 3: A Central Tagging Server vs. a Shared Tagging Application

b. Interface Questions

Each of the different ways that the steve application might be deployed has an effect on the user experience. We need to understand what factors are significant to user's successful engagement with tagging museum objects. For example: How do we help users through the tagging process? How do we motivate them? Is there a way to make tagging fun (Ahn & Dabbish, 2004)? Do we need to guide the process if we wish to encourage more than the free-form assignment of keywords? How will the results change if we lead with facets as a way to guide tagging? Do we need to reward users to get the most out of the tagging experience? How do we adapt interfaces to different needs and executions (mirroring the different uses museums see for tagging and for using the resulting tags)?

The interface experience is a particularly critical issue because variations in it create the plethora of experiences that we hope to support. For example, an institution that is solely looking for additional keywords for objects will need a different interface than one that's looking to offer their visitors sharable galleries of objects based on user tagging patterns. We believe that we can maintain a consistent back-end database application that allows us to change the front-end experience and interaction without sacrificing a common repository for tags and a shared analytical environment. While we currently think this is a workable approach, further testing is needed, not only on the interfaces themselves but also on the very assumption that swappable interfaces will work without compromising the potential of any given user experience.

c. Data Analysis Questions

Architectural decisions have data collection and analysis implications, as different choices result in separate or common data sets that offer varying possibilities for analysis. Single-institution implementations complicate inter-institutional data analysis. We need to understand how to best manage the relationships between tags and museum resources (both to create an experimental data set for analysis and to produce data that is useful to support museum functions.). Ultimately, we need to determine how folksonomic data can be incorporated into museum systems of all kinds, so that we can leverage it to improve visitors' experiences.

d. User Experience Questions

All museums have a common goal: to encourage user engagement with collections. We share a mission to promote and create connections with art. Tagging offers a way to encourage engagement. Folksonomy provides a means to connect on visitors’ terms rather than institutional ones. We’re conscious of the social affordances of our experiment as we move ahead.

But how does the desired experience begin to change the implementation of steve? While one museum may simply see tagging as a method of supporting a better form of search, another may want to create open-ended data sets that visitors can reuse and repurpose in unimagined ways. Yet another may see a venue for visitor studies and market research to inform new kinds of content and experiences. These goals, while not necessarily exclusive, may require different implementations for each experience to feel uniquely good, rather than comprehensively bland. Our challenge becomes to develop steve so that it remains open-ended, without unnecessarily constraining the development of the front and back ends or adding a lot of overhead. As we begin to implement the prototype in multiple contexts we will be watching for interaction effects.

C. Tools and System Architecture

Answering the questions raised in the shared research agenda will require the development of a suite of tools for collecting terms, analyzing data, and monitoring user activities. The Summer 2005 working group at Grindstone Island focused on defining a system that could support the development of a tagging community and enable research on the use of the system. Two working groups – one focused on the front-end (interface) questions and the other on the back-end (database) questions – considered system needs in light of questions being formalized by a working group on the research agenda, and developed specifications for tools to be built. Since then, some progress on building versions of these tools has been made. This section describes the data model proposed by the back-end working group and the interface environments proposed by the front-end working group. The next section of this paper describes work done to date on developing prototypes based on the specifications published by the Grindstone working groups (see <http://www.steve.museum/>).

1. The Data Model

The steve development strategy rests on the assumption that a shared “back-end” database could gather tag data for any collaborating museum. Initially, it appeared that an exquisitely simple data model could support tagging museum objects: an object identifier, an image name, and a place to record the tag text.



Figure 4: A simple model of tagging data

The model is so simple that staff from The Cleveland Museum of Art went home after the “Cataloguing by Crowd” forum and launched their “Help others find me” prototype in a matter of days. Published with little fanfare, the Cleveland “metadating” project has been running since April, 2005. By mid-February, 2006, it had gathered fewer than 250 entries. Many of these entries contained rambling, narrative-style texts (rather than single-term keywords). It was interesting to see that contributors to the tool had discovered a way to engage with the museum in a way that appealed to them, but disappointing that the very simple model for data collection yielded so few usable terms.

if your suggestions are not wholly in English.' At the bottom, there are input fields for 'Your name:' and 'Your email:', and a 'Send Suggestions' button. A thank you message 'Thank you very much for your help.' is at the bottom right."/>

CMA - Help others find Seated Figure

Help others find this object

Seated Figure I 1954.857
Mexico, Oaxaca, Zapotec style, 300 BC-AD 700

Sometimes it can be really hard to find the just the art object you are looking for... You know it had a barn, a woman with a pearl earring, bright bold red stripes or some other feature that really struck a chord. Well, we can't think of every way you might remember a piece, so we need your help.

For the art object at right, please use the box below to type in words and phrases that would be meaningful to you if you were searching for it. Words and phrases might include things you see, emotions you feel, things that this reminds you of, and any other way that in which this object brings meaning to you or makes it memorable.

Enter your suggestions below. Please separate each word or phrase with a semi-colon; this will assure that we know exactly what you mean. Example (not specifically about this object): cow; plow; lace collar; sad; joyful; family; summer. You get the idea.

man; sitting; brown; hat; cross-legged; statue;

We welcome contributions in English or other languages that will help make our collection more widely accessible. Please check this box ☐ if your suggestions are not wholly in English.

Your name:

Your email:

Thank you very much for your help.

Figure 5: “Help others find me” at The Cleveland Museum of Art

Results obtained through the Cleveland implementation were both promising and discouraging, but ultimately provided important feedback that informed the planning by the Grindstone group. In addition, by the time we began discussion in earnest about the data model and term-collection environment, it had become clear that we needed tools that would allow us not simply to collect, manage, and deploy keyword terms, but also to analyze the terms collected and explore the behaviours of our contributors.

The data model issued from the Grindstone meeting (Cataloguing by Crowd Working Group & Hiwiler, 2005b) is significantly more complex than the simple relationship between user, tag, and data shown in Figure 4. This complexity – reflecting the social context of tagging – (shown in Figure 6) distinguishes Grindstone from a number of tagging projects already in place. At Grindstone, we agreed to produce an application that could not only display images and capture corresponding keywords, but also help with the stated institutional goals of engaging communities of users. We wanted a tool that would motivate users to tag, guide them through the process, and reward them when they were done. We wanted “stickiness” – enticing users to spend time in the system and to return for more sessions. We also realized the need for a lot more information about what would make users tag (and tag well) if we were going to create tools that supported and enhanced this motivation. The challenge of recording the contextual data required to analyse the user experience proved too big for a basic add-on to a public Web site.

The back-end group, working in concert with the front-end team, considered the need to manage user sessions, to allow some user choice in selection of images or interfaces, and to manage the presentation of images and data about the work of art. In addition, the group reviewed the developing research agenda, writing a data specification that understood the likely interest in analysing user information according to demographics, object types, user behaviours (such as number of return visits, or length of session).

The back-end group took all of the thoughts elaborated in the research agenda and, in discussion with the front-end group, began to sketch out a model that could support the complex set of demands. The group agreed that modelling the complexity at the outset would strengthen our application, even if we didn't implement all options initially. We envisioned a modular architecture where it was possible to develop and test environments that explored particular variables, all the while gathering tagging data in a standardized form for analysis and use.

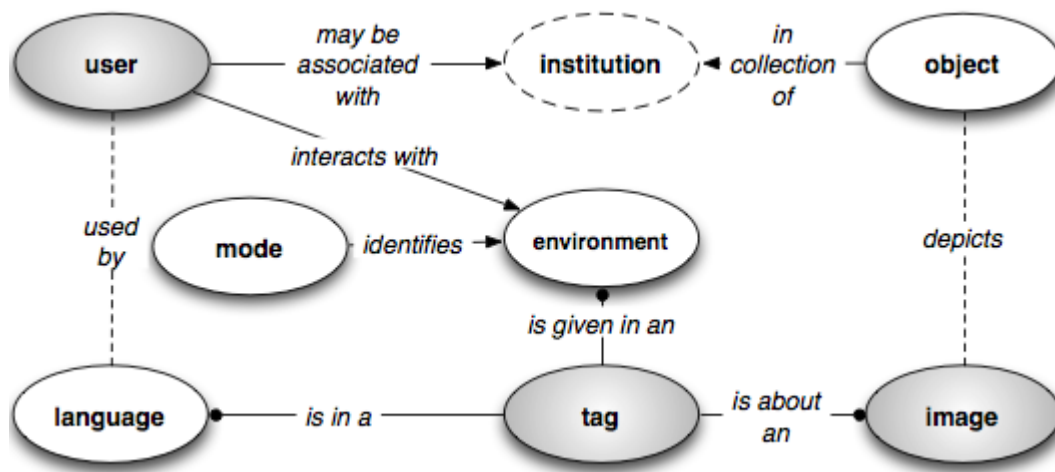


Figure 6: A simplified rendering of the Steve data model

2. Interface Needs

The front-end working group described a fairly complex interface environment, where users could select images they wanted to tag and could choose the layout of the interface from two or more preset designs. Users should also be able to leave, return, and restart a session where they left off. Further, images needed to be filtered so that users would never be presented with repeat images to tag (Cataloguing by Crowd Working Group, 2005). The back-end team envisioned multiple entities and complex relationships. They considered the appeal for the user (demonstrated in the data collected in the Cleveland “metadating” project) of being able to comment on the tagging work and sought ways to capture this user-supplied information while keeping it separated from basic tags. They suggested several possibilities for adding blogging features like the Victoria & Albert Museum’s “Every Object Has A Story” (<http://www.everyobject.net>). They theorized that making the blog function live, so that users could see a blog evolving while tagging, would help to provide additional “stickiness” for the tool.

3. Tool Development Process

Creating tools to support the term collection, user experience, and data analysis functions is expected to be an iterative process. The prototypes we have today are necessarily slim. This has resulted partly because the initial prototype work is being done inside institutional projects that don’t require the complete feature set, or because additional work must be supported through contributions from participating museums or by interested third-party developers such as Think Design, Inc. (<http://www.thinkdesign.com/>).

We expect to engage in this iterative model of development with each successive prototype, adding functionality along the way until our tools reach the full specification. At the same time, we expect to identify flaws in our data model and to modify it as we learn more during the prototype process. Finally, we fully expect that some institutions will be satisfied with a limited subset of the specified features and that they will succeed in producing quality usable data with a basic set of tools. The data model may be complex, but it is also flexible; it will support the feature-limited interface as well as the feature-rich one.

D. Prototypes: our experimental environment

One of the challenges raised in discussions surrounding community cataloguing is how little data we have about how variations in the interface affect the results of having users catalogue. To date, two prototypes have been built based on the specifications developed by the Cataloguing by Crowd working groups. The participants intend to refine and test the second prototype, the “gray prototype,” for a formal launch in Fall, 2006.

1. The Green Prototype

The first beta interface – the green prototype – was developed by the Guggenheim Museum and Tom Leonhardt (tomtom interactive) in advance of the Grindstone Island meeting. Though basic, the prototype served its intended function of supporting discussion of the group’s expectation for a tool to collect, manage, and analyze terms.

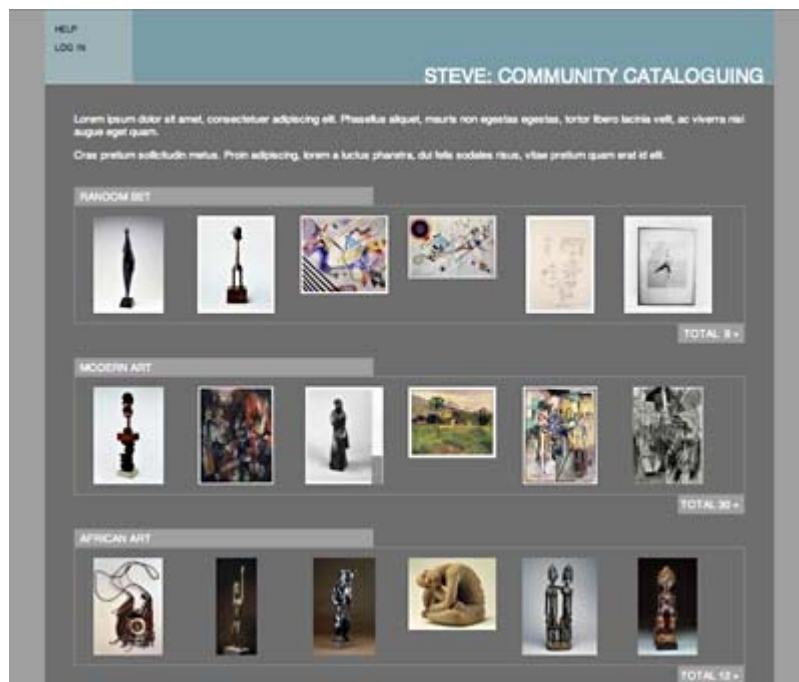


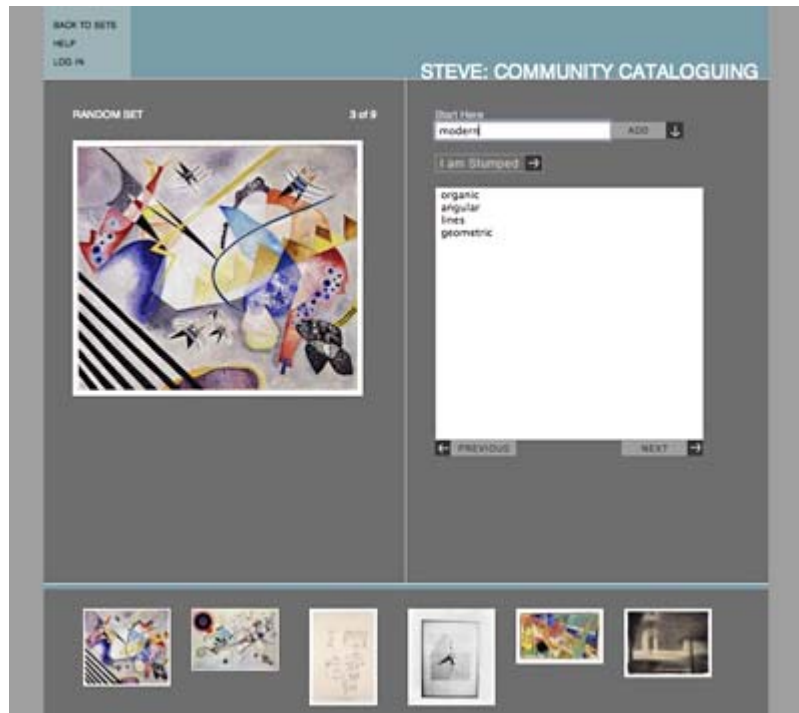
Figure 7: A screen shot from the first “green” steve prototype

The prototype was developed and released under an open source license using the freely available and ubiquitous MySQL/PHP platform. The selected platform is free, making it accessible to any organization with the technical ability to deploy it. It can easily be run on a variety of platforms, including Linux, Mac OSX, and Windows. The software has been released under the GNU Lesser General Public License (LGPL) and code developed through for this project is posted for community use at <http://sourceforge.net/projects/steve-museum/>. By opening access to the source code, we can allow any interested museum or individual to see the code, use it, and ideally contribute to the overall effort by submitting modifications back to the project. The LGPL terms also allow commercial software vendors to use the code with their proprietary products.

2. The Gray Prototype

For the second prototype the Guggenheim and The Metropolitan Museum partnered with Think Design and, using MySQL 4.0 and PHP 4.0, created a system that may be accessed by anonymous visitors or by registered users using the database schema as specified in Steve Data Model (Cataloguing by Crowd Working Group & Hiwiler, 2005b). This environment uses a general SQL handling layer that can support a range of relational databases (portability has been identified as an issue in deployment). The published Web interface is based on a design by Sam Takahashi of The Metropolitan Museum of Art; an XHTML/CSS2 front end architecture allows for other interfaces to be added for testing multiple user environments using the same backend system.





Figures 8 and 9: Screenshots from the second “gray” prototype

a. Basic Functionality

The gray prototype explores a number of functions that put more control of the tagging experience into the hands of the user. Our preliminary tests showed that the tagging experience was significantly different when the groups of objects presented to be tagged were related somehow (thematically, or by style, period, or culture) rather than random. This prototype presents users with a series of image sets from which they can choose. The main cataloguing screen features a vertical term-collection window populated by a single-line text entry field. Tags are added to the collection window by clicking the “add” button or by hitting the return key. Users navigate to the next image using the “next” button. Terms entered by users are recorded and linked to the user, a system time, and the environment in which they were created.

The initial screen allows a user to log in, to register, or to begin cataloguing as an anonymous user. Anonymous users are recorded against a single anonymous user record in the user table, but the environment is identified and recorded on an individual anonymous-user basis. For example, the users table can record multiple anonymous users each with their respective tags associated with them. Later versions may associate users by institution.

b. Populating The Prototype: Image ingestion

Having decided to build a centralized tool for presenting images for tagging and collecting terms, we need to move data from museum systems to the steve tagging tool. Images are uploaded to the site along with the corresponding image data file. The image data file (in XML format) is parsed and processed by a PHP script that loads the data into a table. Each new record is then associated with the path and file name of each image. Currently this is done ad hoc from the command line. Future implementations will allow image ingestion from a Web interface or ingestion from other systems.

c. Administration

At present, the Gray Prototype is only configurable by manually changing configuration files. File and User administration is done from the command line and through directly accessing the SQL tables. Resulting terms are also exported for analysis through direct interaction with the SQL tables.

3. Hooks

The minimum requirements for capturing tag data, essentially an institutional object identifier and a corresponding image, provide opportunities to build “hooks” into existing systems that may already aggregate image and object data. This kind of system-to-system integration may remove the need to transfer data, a significant challenge for

some institutions. One example of an institution's thinking about how hooks to existing local systems might function is The Cleveland Museum of Art's research into integration with its Luna Insight® implementation.

Cleveland is looking at how it might leverage a local implementation of Luna Insight® (<http://www.lunaimaging.com>) to feed steve with sets of images and object data. This source in particular is interesting from the standpoint of developing additional interfaces to steve. For instance, data in Luna Insight® is processed through complex indexing and keyword extraction routines that can be accessed outside of the Luna interface. An internally-developed version of steve might hook into Insight® and, instead of developing a complex ingest routine to move data into the steve tables, might add a search layer that leverages the indexed metadata already present in Insight® to build image sets. This would provide us with another means to test whether tagging behaviour is influenced by the type of content presented for tagging.

A steve ingest model for Insight® content also possible. There are several options for developing image groups within Insight®, then building a reusable ingest routine to harvest those groups for use in steve. For those who are familiar with the user-defined group creation feature in Luna, it may seem a logical source for building image sets for steve. Some work may be necessary to enable this functionality, however, because group elements in Luna Insight® are not stored in an accessible manner; they are stored as hashed text files containing the pointers to the images and are read from within the Insight® software to render the groups on the fly. That said, images within the groups can be selected and exported to produce folders of images and XML files with the corresponding object data. Ingest systems can then be built to extract the appropriate data for steve. Cleveland is currently in discussion with Luna to see if it is possible to directly process the group text files, get the image information from them, and thus eliminate the export step (necessarily a human action), to gather the same information.

Other potential hooks to Luna Insight® that may be worth investigating are Virtual Collections and their XML gateway product. Additionally, the image group feature in Insight® could also be leveraged to produce a reusable export for contribution to a collaborative steve project.

The Open Archives Initiative Protocol for Metadata could also be used by the Museum community to make their images and metadata available for harvesting. The CMA has begun discussions with Cognitive Applications (<http://www.cogapp.com>), about how OAI might be used to support several Cleveland initiatives. If this becomes a more widespread trend and multiple museums adopt the standard, it opens up the possibility to build ingest routines to steve from harvestable content – a significant step forward over moving data manually from system to system, and further facilitating multi-institutional implementations.

E. Next Steps: Functionality to Support Research and Use

While our longer-term systems strategy will evolve based upon iterative testing and development, we recognize the need for the following enhancements to the prototypes in the near-term, to address first-order research questions.

1. User Interface

To test the relationship between interface, user group, and resulting terms, we need to explore additional interfaces and functionality, outlined “A Matrix of Environmental Variables to be Tested in a Museum-Based Folksonomic Tool” (Cataloguing by Crowd Working Group & Bearman, 2005). The following functionality was identified as likely to influence the quality of folksonomic terminology, and therefore to be explored first.

- Zoomable images (to allow examination of images)
- Varying the number of image thumbnails (to focus attention)
- Museum-made groups to tag (pre-defined, related)
- User created group to tag (search-based, personal)
- Users pre-selection of images to tag (another personalization option)
- Essay answer (long text might surface different content)
- Category prompt (suggesting categories might improve consistency)
- User defined category entry (users might identify axes of interest)

2. User Motivation

Additional research needs to be done into ways to make the interfaces “sticky” and to see what types of interface will encourage users to return and continue to tag items.

- Past session return (users might find this more “sticky”)

- A game-based interface or an interface that allows interaction between users and/or museum staff might be desirable.

3. Research Facilitation

To support the creation of environments within which to test particular hypotheses about tagging behaviour, we envision:

- A Web-based feature administration module allowing a single multipurpose interface (allowing all of the optional features to be turned on and off based on an institution's research needs) should be developed.
- Web based export and analysis tools

4. System Administration

To ease the deployment of the steve application, we've identified a number of basic functions:

- A simple Web-based image ingestion method (with a formally published XML import format) is needed to make the application easy to use and ensure that users and advocates are not restricted to users in the technology departments of interested institutions.
- A multi-institution backend should be developed to allow a group of organizations to share a common hosting environment and thus share/reduce the cost and technical acumen required to start using the system.
- Web based features for user administration are needed. Users should have the ability to edit their own profiles and review their tagging results.
- An automated installation routine with a single source file and written documentation needs to be created to ensure widespread adoption.

F. Integrating Tag Data Into Museum Collections

Our first interest in gathering folksonomic data is to determine whether these terms add value. We will analyze whether the terms that have been assigned are related to the content they tag, new to the museum, different from terms we could obtain by processing museum-assigned terminology through existing standard thesauri such as the AAT or ULAN, and if they occur in user queries in our search logs. Several distinct methodological processes are involved in this analysis. "Relatedness" is a subjective judgement that must be applied by people; the best methods for making this kind of judgement need to be developed. Novelty to the museum is a multi-source term-matching process; one of the important questions will be whether these terms can be found in some museum documentation, but not museum documentation that is currently available to users on-line. The existence of terms in folksonomies that might be thesaurally-related to terms the museum has is again a character-string matching issue, though further work would be required to determine whether such terms could reasonably be added in an algorithmic way even if it is found they are present in professional thesauri. Finally, we will ask if the tags that users assign to works contain words that are also found in the search logs of the museum, and of such collective resources as the public catalogue of The AMICO Library.

We anticipate that this analysis will reveal not just lists of terms that are not currently in museum systems, but also the characteristics of such terms. We hope in time to be able to conduct further analysis of the nature of the language that users assign, using linguistic tools such as WordNet. Understanding what aspects of works users describe that are not currently being described by museums could help us to develop data models for description of works of art that satisfy larger numbers of our potential users.

Ultimately, impact on future users and uses depends, of course, on the reintegration of folksonomic terminology obtained through social tagging into the on-line resources of museums. Assuming our findings support the hypothesis that tags add value, we will need to make their details known to our professional colleagues within museums. There are a variety of models of how the additional terms might be deployed in museum Web resources. Obviously it is possible to simply add them to the keywords currently used in museums' object descriptions; but it is likely that the move to add data contributed by the public will lead most museums to identify content authored by the museum and distinguish it from that from public sources. This is probably a healthy and overdue development in any case, as it opens the way for other authorities to comment on works. The same functionality could enable communities with special relationships to works (such as native communities for indigenous art) to juxtapose their understanding with that of the curator, and, of course, to add folksonomy access points without confusing users as to their "authority".

In the coming year, steve will develop a number of post-processing tools to enable these kinds of term analysis, and to support the re-integration of folksonomy into museum documentation systems.

Conclusion

Since the Professional Forum at Museums and the Web 2005, we've made significant progress defining the technical requirements for museum folksonomy applications and creating some prototype tools that set the stage for a serious study of the role of social tagging and folksonomy in the museum environment.

In the coming year we plan to continue testing the interfaces we've designed, to create additional analytic tools for assessing folksonomic terminology and to start exploring the implications of integrating folksonomy into museum on-line resources. We also hope to have proved our hypotheses that social tagging is a fun way to engage communities with collections, and to have shown one way to use Web 2.0 strategies to create museum 2.0.

Acknowledgements

A part of this text draws on a steve proposal for the Tagging Workshop at WWW2006 in Edinburgh, May 2006.

All of the members of the steve team (named at <http://www.steve.museum>) contributed substantively to the ideas presented here.

References

- Ahn, L. v., & Dabbish, L. (2004). Labeling Images with a Computer Game. Paper presented at the CHI 2004, Vienna, Austria.
- Boas, F. (1907). Some Principles of Museum Administration. *Science*, 25(650), 921-933.
- Canadian Heritage Information Network (CHIN). ([2005]). CHIN Data Dictionary: Humanities. Retrieved November 9, 2005, from <http://daryl.chin.gc.ca:8000/BASIS/chindd/user/wwwhe/SF>
- Cataloguing by Crowd Working Group. (2005). Cyber-Cataloguer: Front End Functionality. Retrieved August 25, 2005, from <http://www.steve.museum/docs/cyberCataloguerFR.html>
- Cataloguing by Crowd Working Group, & D. Bearman. (2005). A Matrix of Environmental Variables to be Tested in a Museum-Based Folksonomic Tool. Retrieved 2005, Aug 24, from <http://www.steve.museum/reference/CCA-environmentMatrix20050.pdf>
- Cataloguing by Crowd Working Group, & D. Hiwiler. (2005a). "Back End" Working Group Report -2. Retrieved August 25, 2005, from <http://www.steve.museum/reference/ComCatBackEnd050722.pdf>
- Cataloguing by Crowd Working Group, & D. Hiwiler. (2005b). Steve data model. Retrieved August 25, 2005, from <http://www.steve.museum/static/reference/SteveDataModel20050728.pdf>
- Cataloguing by Crowd Working Group, & M. Jenkins. (2005). "Back End" Working Group Report -1. Retrieved August 25, 2005, from <http://www.steve.museum/reference/ComCatBackEnd050720.pdf>
- Cataloguing by Crowd Working Group, & J. Trant. (2005). Research Questions: A Summary of Issues Raised in Discussion. Retrieved August 25, 2005, from <http://www.steve.museum/reference/ComCatResearchQ050723.pdf>
- Chudnov, D. (2005). visualizing folksonomy data from unalog using Starlight. One Big Library [blog] Retrieved January 10, 2006, from <http://onebiglibrary.net/node/13>
- Chun, S., & M. Jenkins. (2005). Cataloguing by Crowd; a proposal for the development of a community cataloguing tool to capture subject information for images (A Professional Forum). Museums and the Web 2005. Retrieved August 7, 2005, 2005, from http://www.archimuse.com/mw2005/abstracts/prg_280000899.html
- Cleveland Museum of Art. (2005). Help others find this object. Cleveland, Ohio. [linked to all objects in their collections search at <http://www.clevelandmuseumofart.org/Explore/>].
- Durbin, G. (2004). Learning from Amazon and eBay: User-generated Material for Museum Web Sites. Museums and the Web 2004: Proceedings. Retrieved August 7, 2005, 2005, from <http://www.archimuse.com/mw2004/papers/durbin/durbin.html>

- Golder, S. A., & B. A. Huberman. (2005). The Structure of Collaborative Tagging Systems Journal of Information Science (preprint) <http://www.hpl.hp.com/research/idl/papers/tags/tags.pdf>
- McKenna, G., & E. Patsatzi. (2005). *SPECTRUM: The UK Museum Documentation Standard* (Version 3.0 ed.). Cambridge: MDA.
- Millen, D., J. Feinberg, & B. Kerr. (2005). Social Bookmarking in the Enterprise. *Social Computing*, 3(9). http://acmqueue.com/modules.php?name=Content&pa=printer_friendly&pid=344&page=1
- Powerhouse Museum, & S. Chan. (2005, ongoing). Electronic Swatchbook. Retrieved January 10, 2006, from <http://www.powerhousemuseum.com/electronicwatchbook/>
- Vander Wal, T. (2005). Delicious Lesson and Social Network Ecosystems. Off the Top [blog] Retrieved December 22, 2005, from <http://www.vanderwal.net/random/entrysel.php?blog=1765>
- Walsh, P. (1997). The Web and the Unassailable Voice. *Archives and Museum Informatics*, 11(2), 77-85. Available <http://www.archimuse.com/mw97/speak/walsh.htm>
- Windisch, N. K. (2005 -). flickrTagFight. Retrieved March 6, 2006, from <http://www.netomer.de/flickrtagfight/fight>

Cite as:

Wyman, B., et al., Steve.museum: An Ongoing Experiment in Social Tagging, Folksonomy, and Museums, in J. Trant and D. Bearman (eds.). *Museums and the Web 2006: Proceedings*, Toronto: Archives & Museum Informatics, published March 1, 2006 at <http://www.archimuse.com/mw2006/papers/wyman/wyman.html>

last updated:

March 7, 2006 10:47 AM

Archives & Museum Informatics, 158 Lee Avenue, Toronto, Ontario, M4E 2P3 Canada

Telephone: +1 416 691 2516 | Fax: +1 416 352 6025 | E-mail: [mw2006 @ archimuse.com](mailto:mw2006@archimuse.com)

Copyright © 2006 - Archives & Museum Informatics - All rights reserved.