

SOCIAL CLASSIFICATION AND FOLKSONOMY IN ART MUSEUMS: EARLY DATA FROM THE STEVE.MUSEUM TAGGER PROTOTYPE

Jennifer Trant <jtrant@archimuse.com>

Archives & Museum Informatics

and

Faculty of Information Studies

158 Lee Avenue

University of Toronto

Toronto, Ontario M4E 2P3, Canada

Toronto, Ontario M5S 3G6, Canada

Abstract

The collections of art museums have been assembled over hundreds of years and described, organized and classified according to traditions of art historical research and discourse. Art museums, in their role as curators and interpreters of the cultural record, have developed standards for the description of works of art (such as the *Categories for the Description of Works of Art*, CDWA) that emphasize the physical nature of art as artefact, the authorial role of the creator, the temporal and cultural context of creation and ownership, and the scholarly significance of the work over time. Collections managers have recorded conservation, exhibition, loan and publication history, along with significant volumes of internal documentation of acquisition and storage, that support the custody and care of artefacts of significant cultural value. But the systems of documentation and classification that support the professional discourse of art history and the management of museum collections have failed to represent the interests, perspectives or passions of those who visit [use?] museum collections, both on-site and online. As museums move to reflect the breadth of their audiences and the diversity of their perspectives, so must museum documentation change to reflect concerns other than the traditionally art historical and museological.

Social tagging offers a direct way for museums to learn what museum-goers see in works of art, what they judge as significant and where they find or make meaning. Within the steve collaboration(<http://www.steve.museum>), a group of art museums is collectively exploring the role of social tagging and studying the resulting folksonomy (Bearman & Trant, 2005; Chun, Cherry, Hiwiller, Trant, & Wyman, 2006; Trant & Wyman, 2006). Analysis of terms collected in the prototype steve tagger suggests that social tagging of art museum objects can in fact augment museum documentation with unique access points not found in traditional cataloguing. Terms collected through social tagging tools are being compared to museum documentation, to establish

the actual contributions made by naïve users to the accessibility of art museum collections and to see if social classification provides a way to bridge the semantic gap between art historians and art museums' publics.

1. Introduction

Locating and gaining access to the primary sources of art history – the works of art themselves – is one of the major challenges of a student, scholar, or enthusiast engaged with modes of visual expression. Works of art related by subject, theme, artist, or other scholarly interest can be found in public and private collections around the world. Often, much of the 'work' in art historical scholarship is in identifying works appropriate for study, and building personal collections of textual and visual documentation to support research. Indeed, departments of art history – and many museums – have large Visual Resources collections to satisfy just this requirement.

Traditions of sharing information about cultural collections through the creation and distribution of reproductions stretch back to the creation of reproductive prints depicting famous paintings in the Renaissance, and were formalized in international treaties supporting the creation and exchange of plaster casts in the 19th century. But the possibility of uniting information about dispersed cultural collections has most actively engaged those responsible for managing documentation about art collections with the development of networked communications. "Virtual databases", logical constructs that bring together information resources housed in distinct databases and maintained by different institutions, began to seem possible when the Internet offered ways to connect disparate text database resources. But it was the World Wide Web, which facilitated the development of more approachable interfaces and enabled the easy integration of text and image resources that accelerated development. Museums have moved (in the last ten years) from wondering whether they should put their collections on-line to exploring the implications of having their collections on-line. This openness has coincided with an increased focus on the role of museums in the community, and with the development of more user-centered philosophies for the creation and delivery of networked information resources, and is chronicled by the Proceedings of the Museums and the Web conferences chart this change. (see Bearman & Trant, 1997; through Trant & Bearman, 2006).

When viewed from a user perspective, on-line museums collections, while a vast improvement on the limited access offered previously, may not be fully satisfactory. Networked information

resources still mirror physical museum reality in many ways. It is still not possible to search art museum collections as a whole; one must separately visit each museum site. The information presented is structured according to museum goals and objectives – which may not mesh with those of the user. The language used is often highly specialized and technical, rendering resources inaccessible or incomprehensible. A work of art or other museum object may be embedded in an exhibition or other interpretive context with a point-of-view not be shared by the user. Or inversely, it may only appear in a database, completely de-contextualized and without the meaning that comes from seeing it alongside other artefacts of the same culture.

The challenge of creating and organizing personal collections of networked information resources is not unique to the users of on-line art museum information. Others who actively use Web-based resources, such as online databases of scientific articles, have begun to develop tools to enable the creation of personal collections of ‘bookmarks’ or pointers to networked resources, that are described or ‘tagged’ with words that identify and describe them (Mathes, 2004; Quintarelli, 2005); see del.icio.us (Golder & Huberman, 2005), connotea (Hammond, Hannay, Lund, & Scott, 2005; Lund, Hammond, Flack, & Hannay, 2005) and PenTags (Pennsylvania State University Library, 2005-) for representative examples. “Social Tagging” refers to the practice of publicly labeling or categorizing resources in such a shared environment. The resulting assemblage of tags form a “folksonomy”: a conflation of the worlds ‘folk’ and ‘taxonomy’ used to refer to an informal, and organic assemblage of related terminology (Vander Wal, 2005). When shared with others, or viewed in the context of what others have tagged, these collections of resource identifiers, tags and people begin to take on additional value through network effects. Searching tags supports the discovery of relevant resources, and the social relationships that develop between taggers themselves become a means of information discovery (Marlow, Naaman, Boyd, & Davis, 2006).

This paper examines the access offered to art museum collections on-line and reports on early results from a prototype tagging application, the steve tagger that explores whether social tagging and the resulting folksonomy could improve that access. Social tagging and folksonomy should be considered within the context of other strategies for improving access to art information. Each strategy offers some benefits, and many are likely to be used in concert to facilitate access to art museum collections.

The content of art museum collections is visual, but we work with the ideas represented in them in a primarily textual mode. This produces the major paradox in the documentation and retrieval of art museum collections. What is searched is not the work of art, or even a reproduction (however faithful) of the work of art, but a textual representation of those characteristics of the work of art that were seen as salient by its custodian and/or descriptor. On-line, access may be improved, but it is limited by the nature of available searchable metadata describing a work of art, or by the capabilities of image processing when queries are made by image (visual) content.

In contrast to searching the free-text of digitized books, for example, which provides a direct form of access to their content, what we are able to search in digitized art museum collections is a limited re-presentation of their content, transformed into another media. What is retrieved is not the original or even a facsimile (with many functional equivalents to the original that may be considered to carry the same information content): what is retrieved is a surrogate (Roberts, 1994). It may only be a structured textual description, or it may be a structured textual description accompanied by other information types (image, sound, multimedia, prose). None of these representations can be said to completely stand for, or be considered functionally equivalent to the original work of art; but all serve to enhance access to dispersed works of art, and ease some of the basic work of art history.

2. Finding Works of Art On-line

Museum collections management systems core information found in the majority of on-line collections catalogues. The challenges of adapting this content to fulfill needs of public access are many (Kydd, MacKenzie, & Myles, 1998). The access offered by this kind of information is best illustrated by example.

2.1 A sample query and work description: Museum of Fine Arts, Boston

At the Museum of Fine Arts, Boston, where over 320,000 descriptions of works of art are searchable in a Web-accessible database (Figure 1) it is possible to search by Accession Number, Object Name/Title, Artist/Maker, Object Place, Medium, Culture, Classification, Credit Line, Provenance, Image, On View and Keyword. Of these categories, two relate to the primary art historical concern for consulting visual evidence: “On View” limits a search to those works

currently installed in the museum galleries. “Image” limits the search to those records that have an accompanying digital image.

The screenshot shows the Museum of Fine Arts, Boston Collections Search Results page. The page has a red navigation bar with links for HOME, VISIT, CALENDAR, EXHIBITIONS, COLLECTIONS, and ABOUT THE MFA. The main content area is titled "Collections Search Results" and features a search form with the following fields: Accession Number, Object Name/Title, Artist/Maker, Medium, Culture, Classification (with a dropdown arrow), Credit Line, Provenance, Image (with a checkbox and the text "(find only objects with images)"), On View (with a checkbox and the text "(find only objects on view)"), and Keyword. Below the form is a "Search" button and a "Sort by:" section with "Title" and "Ascending" dropdowns. To the right of the search form, there is a "Collection Search" box showing "330,321 artworks" and an "Advanced Search" link. A "Please note" section provides information about the collection and image rights. The footer includes accessibility, contact, sitemap, and RSS/XML feed links.

Figure 1. The Museum of Fine Arts, Boston, offers searchable access to an online database of 330,321 objects (as of September 15, 2006). Search fields include Accession Number, Object Name/Title, Artist/Maker, Object Place, Medium, Culture, Classification, Credit Line, Provenance, Image, On View and Keyword. (see http://www.mfa.org/collections/search/search_art.asp).

A search for “Expulsion” in keywords finds a number of works, including Benvenuto di Giovanni's *Expulsion from Paradise* (1470s) (Figure 2). The data accompanying describing this work provides a good illustration of the variety inherent in museum documentation. A seemingly full record (on first glance) provides significant background about “Provenance/Ownership History”, but little else beyond the standard museum label copy/text. Some Classification terms are included – “Religious – Old Testament” and “Nude” – but no further detail is given about the iconography of the work – the Expulsion from Paradise – or of the main protagonists: Adam and Eve, the serpent, or he who did the expelling.

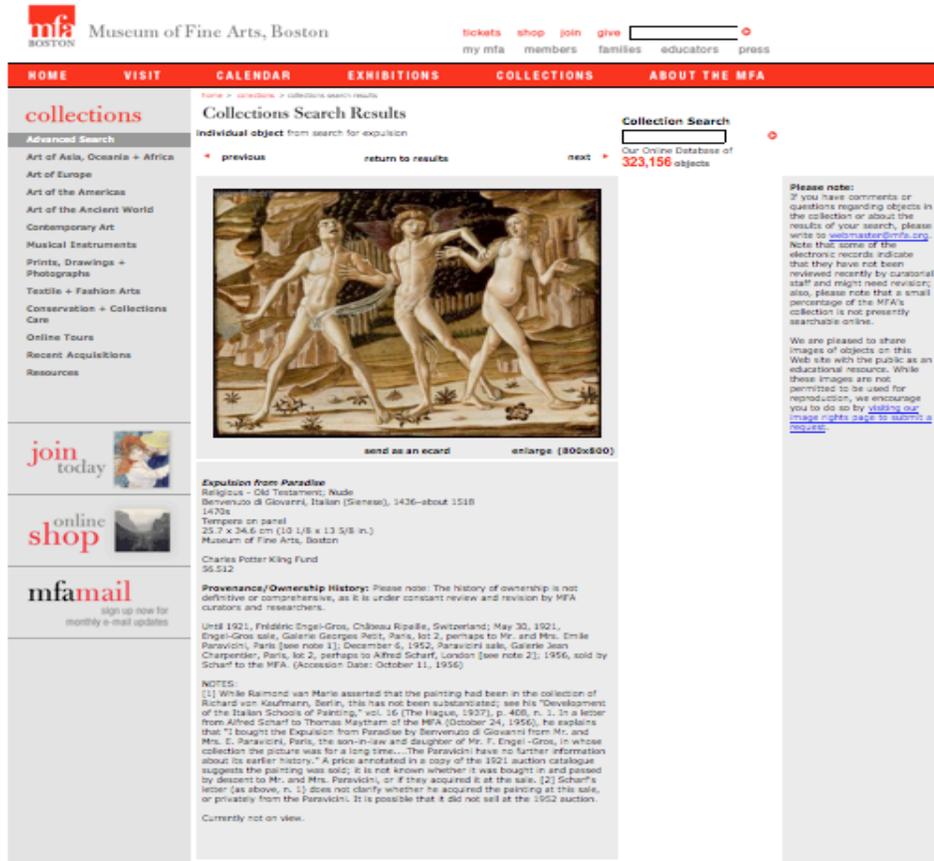


Figure 2. This online record for Benvenuto di Giovanni's *Expulsion from Paradise* in the collection of the Museum of Fine Arts, Boston Collections Database includes a detailed history of the provenance of the painting, including scholarly notes about previous assertions of ownership. Subject classification includes "Religious - Old Testament" and "Nude"; it does not mention man, woman, Adam or Eve.

2.2 A sample query and work description: The Metropolitan Museum of Art

Simple searching of on-line art museum collections is common. Keyword search boxes commonly appear on Web sites, on the Collections pages of the Museum of Fine Arts, Boston (Figure 2) or in the main navigation at The Metropolitan Museum of Art (Figure 3). Where at Boston, the desire to search the collection is imputed from the presence of the collections search box only in that sub-section of the site, at The Metropolitan Museum of Art, the search box is on every page, and the visitor is asked what to search for – by type of content or site section (e.g., Works of Art, The Met Store, Visitor Information).



Figure 3. Searching can be as simple as a keyword box on the home page of a museum web site, as is the case at The Metropolitan Museum of Art. However, here you must specify what you wish to search for Works of Art, The Met Store, Visitor Information

A search for “expulsion” in Works of Art at The Metropolitan Museum of Art retrieves a number of works, including Charles Joseph Natoire’s *The Expulsion from Paradise* Figure 4. As well “label copy” (artist, title, date, materials), this work’s on-line record includes a textual “Description”. This work’s Description reflects the interests of the museum – it’s a “ well preserved painting...” – and the art historian –“when he was living in Paris near his contemporary François Boucher”. But nothing is said about the iconography, beyond that it is typical of “an eighteenth-century inclination towards intimate interpretations of religious themes”.



Figure 4. A screen shot showing on-line documentation of a Work of Art from The Metropolitan Museum of Art's European Painting Collection: Charles Joseph Natoire's *The Expulsion from Paradise*, 1740.

2.3 An example query: CHIN's Artefacts Canada database

Curatorial notes are not always present in on-line museum collections documentation of works of art. For example, the Canadian Heritage Information Network's Artefacts Canada: Humanities contains "almost 4 million object records and 400,000 images from hundreds of museums across the country" {Canadian Heritage Information Network (CHIN), 2006 #2279}. Searching in Artefacts Canada can be by simple keyword, or by more advanced facets, organizing CHIN Humanities Data Dictionary (Canadian Heritage Information Network (CHIN), 2002) fields around "who, what, where, when and how". But works are often described in a very cursory manner. For example, the record for Hortense Gordon's *The Space Time Field* (Figure 5) tells us very little, except that this work is a "non-objective", Canadian abstract painting.

Works of art can be approached and interpreted from many different perspectives; documentation represents one of many different points of view (Sledge, 1995; Sledge & Case, 1995) and, as a result of necessary choices, some things important to users might not be mentioned at all (Honigsbaum, 2005).

These differences in perspective may be exacerbated by ‘traditional’ documentation practice. Theories of indexing and retrieval have been constructed in a climate of scarcity. Precision was essential – and preferable over recall – when pricing structures for the use of research networks were based on a fee per search (as they were in RLG in the 1980s) and on the number of records returned. But end-user searching is now far more common than professionally mediated searches. Browsing larger results is an accepted form of query refinement – one expected in art history, where looking is a professional skill, and becoming more common in image searching generally (Cunningham & Masoodian, 2006).

Recommendations in the literature of image retrieval for improving access to visual collections often focus on establishing effective vocabulary and authority control for the description of visual collections (Harpring, 2002; Shubert, 1996). But even when their expression is rigorously controlled, the concepts represented in the records museums make available may not reflect the interests of museum users. Studies of queries of museums (McCorry & Morrison, 1993) and museum information resources (Janney & Sledge, 1995) begin to reveal gaps between the professional framework of museum documentation and the perspective of users of museum collections. Studies that begin with user needs (such as Elinich, 2004; Reich & Lindgren-Streicher, 2006; Samis, Mitroff, & Johnson, 2005; Schaller & Allison-Bunnell, 2003; Stephenson & McClung, 1998) surface differences between what museums have available and what users expect or want.

Museums are aware of the challenges of putting collections information on-line (Reilly, 2000). Museum information professionals know that people are “searching for meaning, not just records” (Doolan, Peacock, & Ellis, 2004), and strive to provide a number of different ways to encounter collections on-line (including exhibitions, in-depth features, publications, games, and educational materials for teachers as well as collections databases). There is an irony that for some kinds of users, making collections databases available on-line may not make collections themselves more accessible.

4. Potential for Social Classification / Folksonomy

It is within this context that a group of American art museums, and the professionals that support them, have come together in the to explore the potential for social tagging and folksonomy to enable access to art museum collections on-line in a project we call steve (Chun et al., 2006). Inspired by popular social tagging environments like flickr and del.icio.us, and encouraged by the success of the ESP Game (Ahn & Dabbish, 2004), this group has been building an environment within which to research the contribution of publicly assigned terms to the on-line accessibility of art collections (available at <http://www.steve.museum>).

Early proof of concept tests at The Metropolitan Museum of Art (Trant, 2006) revealed striking differences between the terms assigned to works of art by professional art historians and librarians, and those assigned by non-professionals. There was a significant semantic gap, between professional and public discourse about works of art that could be bridged by incorporating user-supplied terms into art museum documentation. Supporting social tagging of art collections, and integrating the resulting folksonomy into on-line art museum collections search, seemed promising enough as an additional access strategy to warrant serious exploration. Proof of Concept studies seemed to indicate that terms assigned by non-specialists could significantly enhance the number and kind of points of access to works of art, and could offer another layer of documentation to supplement and complement that provided by professional cataloguers.

4.1 Preliminary results from the steve tagger prototype

The steve.museum group has developed a prototype environment that allows the tagging of works of art from participants' collections. *Our goal with the prototype is to understand the conditions that will affect social tagging and folksonomy in art museums.* We are now launching a two-year study of the nature of social tagging and folksonomy in art museums that builds on the preliminary results reported here.

The steve tagger that has been available in a number of forms since an alpha version informed discussion at the first working meeting of the project (Cataloging by Crowd Working Group & Leonhardt, 2005). The data discussed here was collected in the second version of the steve tagger, the "gray prototype" between October 2, 2005 and September 28, 2006. During this time, the

tagging tool was available, linked from <http://www.steve.museum>, in an informal beta test environment, shown in Figure 5, Figure 6 and Figure 7.

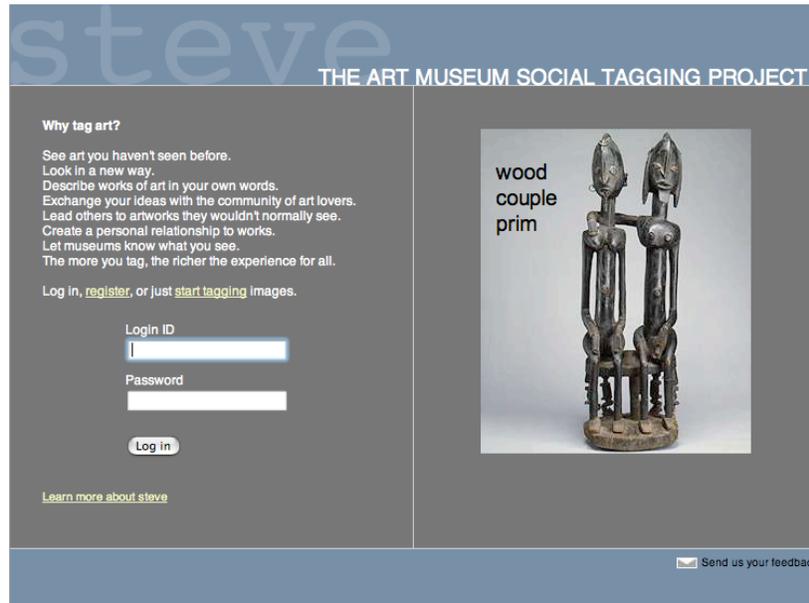


Figure 5: Login screen from the steve.museum tagging tool: gray prototype as made available in 2006. Login is optional. Screenshot taken September 28, 2006.

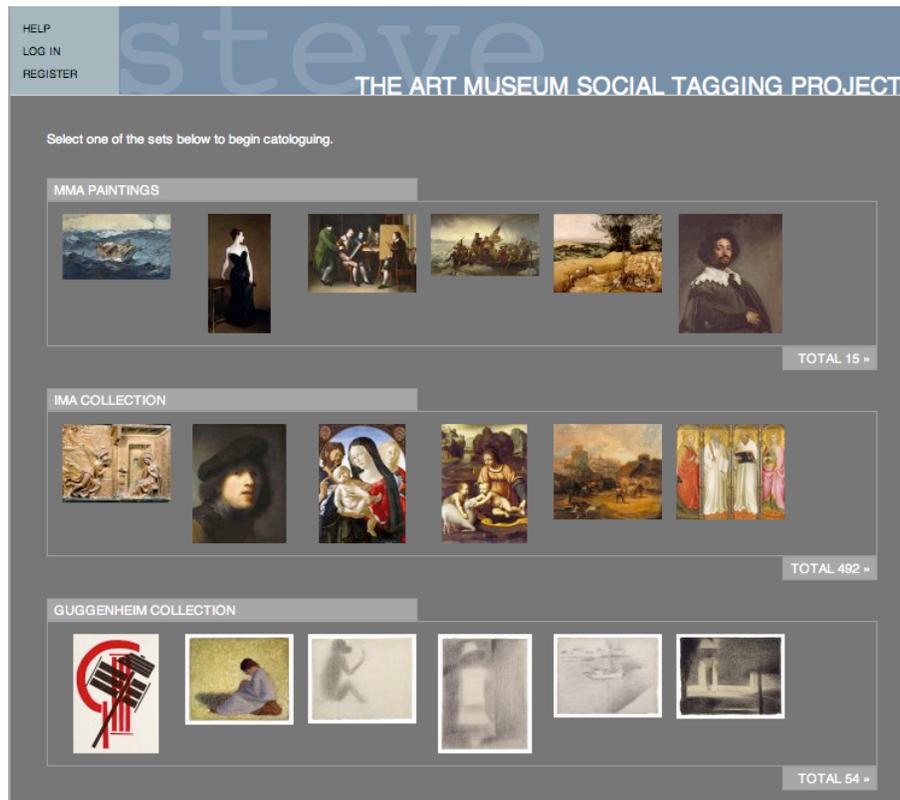


Figure 6. sets screen from the steve.museum tagging tool: gray prototype as made available in 2006. Users are presented with a series of sets of works to tag. Screenshot taken September 28, 2006.

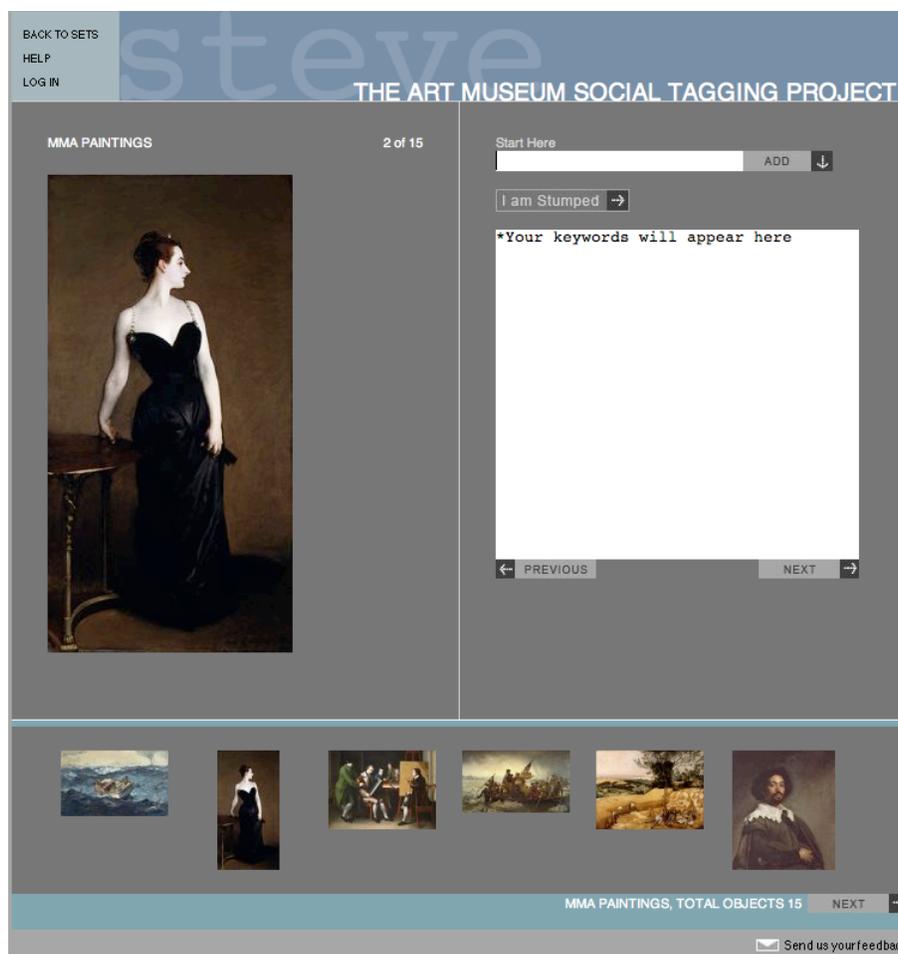


Figure 7: Tagging screen from the steve.museum tagging tool: gray prototype as made available in 2006.

The tagging interface shows a work, and prompts for tags. Other works in the set are shown below.

Screenshot taken September 28, 2006.

4.2 Tags over time

A total of 1,313 works from 4 institutions were added to the steve tagger over the course of the prototype period, October 2005-September 2006. A total of 7,339 tags were assigned to these works (Chart 1). Little recruiting was done to draw people to the tagger, beyond presentations at professional conferences. No conscious attempts were made to draw in public participation: though one mention in the *New York Times Magazine's* Year in Ideas (Pink, 2005) drew popular attention to the project, it didn't publish a site address. As would be expected, tagging activity peaked shortly after each of the steve presentations made in 2006.

17th SIG/CR Classification Research Workshop, November 4, 2006

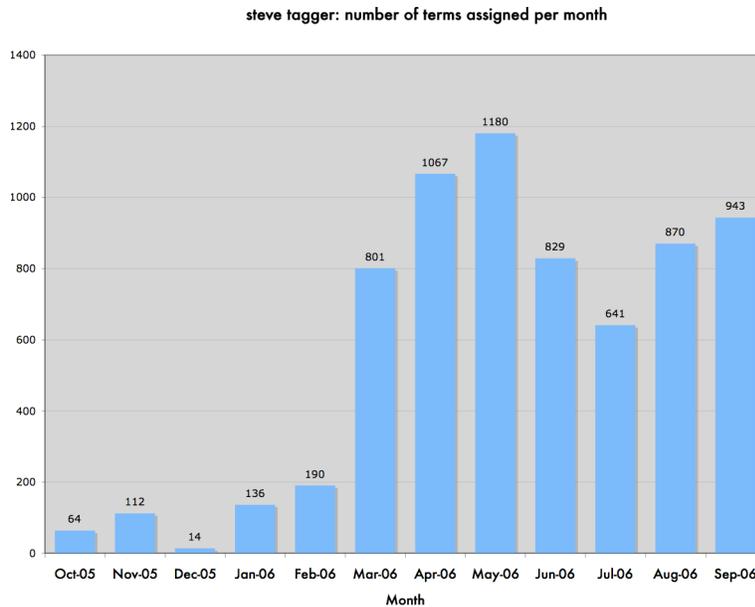


Chart 1: Number of Terms per Month

4.3 Users

As is shown in Figure 5, the log-in screen from the steve tagger, registration was optional. Users could either create an identity at steve.museum that enabled them to build a tagging history over time, or they could just ‘begin tagging’. During the time that the prototype was available 137 users created accounts at steve.museum. This number includes 13 members of staff from steve.museum institutions who created accounts and tagged works as if they were users of the system. *Table 1: Total Number of Terms Assigned by Type of User* shows, the average number of terms assigned by steve.museum participants (23.0) was almost identical to that of other Known Users (22.5). steve.museum participants did assign a higher median number of terms, but did not have as marked highs or lows as other Known Users. The terms assigned by steve.museum participants were not excluded from this experimental data analysis.

| | average | high | low | median |
|--------------------|---------|-------|-----|--------|
| Known Users | 22.5 | 451.0 | 1.0 | 7.0 |
| Steve Participants | 23.0 | 59.0 | 3.0 | 19.0 |

Table 1: Total Number of Terms Assigned by Type of User

4.4 Valid tags

Whenever the possibility of public tagging of works of art is discussed, concerns about the quality of publicly supplied data are raised (Cataloguing by Crowd Working Group & Trant, 2005). Will

tags assigned by the general public contain a large number of errors? Terms might be spelled wrong, or might not be appropriate for the museums’ general audience.

The 7,339 terms collected in the prototype steve tagger were reviewed; entries with spelling mistakes, terms in foreign languages, and data entered that was not ‘words’ were identified. Only 492 of the 7339 terms (6.7%) were discarded. These are summarized in *Table 2: Terms Disqualified as Inappropriate*.

| | Number of Terms | Percentage of Disqualified |
|------------------------------|-----------------|----------------------------|
| Spelling Error | 193 | 39.2% |
| Not Word | 160 | 32.5% |
| Foreign Language | 86 | 17.5% |
| Punctuation or Character Set | 36 | 7.3% |
| Questionable | 13 | 2.6% |
| Obscenity | 4 | 0.8% |
| Total | 492 | 100.0% |

Table 2: Terms Disqualified as Inappropriate

Of the terms that were disqualified, the majority (193 or 39.2%) represented spelling errors. These were words that were not spelled in either American or British English, and included a number of strings such as “blackandwhite” or “chestofdrawers” that were the conditioned result of other tagging tools (like del.icio.us) that do not allow multiple word tags. The next group comprised data entry that was ‘not a word’ (160 or 32.5%); it included a significant number of random character strings, like ‘hfadhf’ entered experimentally to see how the interface behaved. Foreign language terms (86 or 17.5% of the total) were also disqualified from this review, though there are times in the discourse of art history when foreign language terminology is appropriate. A group of terms were removed because of early problems with character sets and punctuation (36 or 7.3%) later corrected in the prototype environment, and a group of questionable terms were excluded, but flagged fur further follow-up. Only four (4 or 0.8%) of the disqualified terms were obscenities, representing 0.05% of the total number of terms entered in the prototype steve tagger.

Table 3: Terms by Type of User shows that slightly more of the invalid terms (346 or 4.7% of all terms) came from users who did not register. Of the 492 terms excluded (6.7% of the total number of terms), Known Users and steve participants supplied 2% and Unknown Users 4.7%; Unknown Users made slightly more than twice as many errors in tagging as known users. A total of 93.3% of terms collected in the steve tagger (6,849 of 7,339) were judged ‘valid’ terms.

| | Number of Users | Good Terms | | Bad Terms | | Total Terms |
|--------------------|-----------------|------------|-------|-----------|------|-------------|
| Known Users | 124 | 2677 | 36.5% | 118 | 1.6% | 2795 |
| steve Participants | 13 | 271 | 3.7% | 28 | 0.4% | 299 |
| Unknown Users | ? | 3899 | 53.1% | 346 | 4.7% | 4245 |
| All Users | 137 | 6847 | 93.3% | 492 | 6.7% | 7339 |

Table 3: Terms by Type of User

This brief analysis represents only the first-phase of review of the value of tags. It does not assess the applicability of the tag to the work, a judgment that requires a work-by-work examination, to determine if the terms supplied are applicable to the work of art, or whether they are inaccurate or misleading. Nor does it examine errors – such as spelling errors or typos – to see if they represent useful points of access.

4.5 Tags and works

It is not possible to draw many conclusions about tagging behaviour from the number of tags assigned to each work in the steve tagger prototype. Works were added to the tagger at different times, as participating institutions became comfortable with the idea of a test environment within which members of the general public would assign terms to works of art from museum collections. Works presented to users were not randomized, so the works that appeared early in the set were tagged a larger number of times than those that appeared later in the set. Chart 2 shows the strong clustering of tags on a few works (those presented ‘early’ in the interface) and makes the case for a tagging interface that consciously varies the works presented to users, in order to ensure a more even distribution of tags.

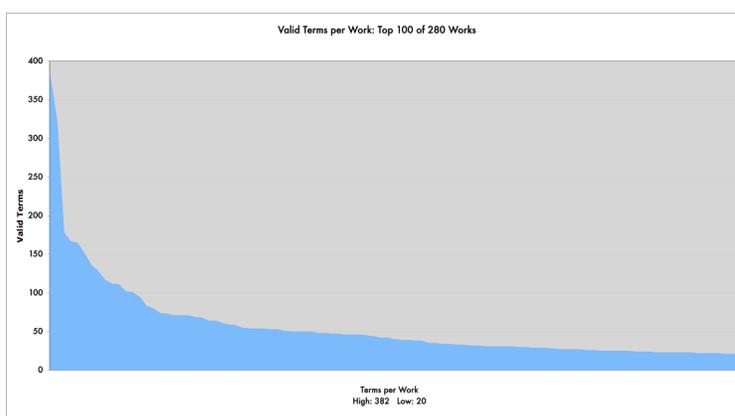


Chart 2: Valid Terms per work for the top 100 of 280 works in the steve tagger.

The tags gathered in the prototype do confirm some of the results of the Proof of Concept studies, and point to directions for further research.

17th SIG/CR Classification Research Workshop, November 4, 2006

Chart 3 shows tagging statistics for the four most-tagged works. The works with the most tags

assigned to them  (382), and  (338) were the first two presented in the first group in the interface as it appeared in September 2006. The other two works in the top four also appear in the first set presented for tagging, on the first row of the ‘sets’ screen. But the number of valid tags assigned to them (178 and 166) represents a major fall-off. There is a corresponding 50% drop in the number of known users tagging the first two works as opposed to the subsequent two works. The disproportionately high number of terms assigned to the first two works in the upper left hand corner of the ‘sets’ screen confirms the propensity for users who are ‘just trying it out’ to click on the first item(s) encountered. Means of maintaining user motivation is a question for further study.



| Ranking: By Number of Terms | 1 | 2 | 3 | 4 |
|--|------|------|-------|-------|
| Position in Interface | 1 | 2 | 4 | 5 |
| Number of Terms Assigned | 390 | 338 | 187 | 171 |
| Invalid Terms (not words) | 8 | 16 | 9 | 5 |
| Valid Terms (words) | 382 | 322 | 178 | 166 |
| Unique Terms | 117 | 132 | 76 | 84 |
| Appropriate Terms | 113 | 130 | 73 | 76 |
| Inappropriate Terms | 4 | 2 | 3 | 8 |
| Number of Known Users | 32 | 29 | 15 | 16 |
| Average Number of Terms per Known User | 4.75 | 4.14 | 5.13 | 4.50 |
| Known Terms | 40 | 54 | 36 | 42 |
| Unique Known Terms | 8 | 10 | 8 | 11 |
| Percentage Known | 6.8% | 7.6% | 10.5% | 13.1% |
| New Terms | 340 | 268 | 142 | 124 |

Chart 3: steve tagger examples. Top four works by number of valid tags, all from *The Metropolitan Museum of Art*, and all in the first set presented for tagging

17th SIG/CR Classification Research Workshop, November 4, 2006

There is a strong cluster of terms at the top of the curve, dropping off fairly quickly into terms assigned only once: 38.9% of the unique terms were assigned more than once; 60.1% of the unique terms were assigned only once. This represents an exceptionally strong agreement in the first group of terms. Inter-tagger consistency was very high (Markey, 1984): 68% of known users assigned the most common term (boat), 56.3% the second most-common term (storm). The nature of these terms – boat, storm, sharks, sea, ocean – shows a lacunae in the scholarly museum documentation. This strong co-occurrence of the most commonly assigned terms holds true across the four most-tagged works (Chart 5: steve tagger: consistency in terms across known users; Top five tags for each of the top four works; all show a very strong clustering around the most commonly assigned terms. The strong agreement on the first group of terms for each work is significant. If this holds true, it points to ways that museums can use statistical tests to validate terms gathered through social tagging.



Winslow Homer (1836–1910)
The Gulf Stream, 1899
Oil on canvas; 28 1/8 x 49 1/8 in.
(71.4 x 124.8 cm)
Catharine Lorillard Wolfe Collection,
Wolfe Fund, 1906 (06.1234)



John Singer Sargent (1856–1925)
Madame X (Madame Pierre Gautreau), 1883–84
Oil on canvas; 82 1/8 x 43 1/4 in.
(208.6 x 109.9 cm)
Arthur Hoppock Hearn Fund, 1916
(16.53)



Emanuel Leutze (1816–1868)
Washington Crossing the Delaware,
1851
Oil on canvas; 149 x 255 in. (378.5 x
647.7 cm)
Gift of John Stewart Kennedy, 1897
(97.34)



Pieter Bruegel the Elder (Netherlandish,
active by 1551, died 1569)
The Harvesters, 1565
Oil on wood; Overall, including added
strips at top, bottom, and right, 46 7/8 x
63 3/4 in. (119 x 162 cm); original painted
surface 45 7/8 x 62 7/8 in. (116.5 x 159.5
cm)
Rogers Fund, 1919 (19.164)

| Known Users | term | | | term | | | term | | | term | | |
|--------------------------|--------|--------|-------|-------------|--------|-------|------------|--------|-------|-----------|--------|-------|
| | term | number | % | term | number | % | term | number | % | term | number | % |
| Consistency in Top Terms | | | | | | | | | | | | |
| 1 | boat | 22 | 68.8% | black dress | 9 | 31.0% | washington | 7 | 46.7% | harvest | 8 | 47.1% |
| 2 | storm | 18 | 56.3% | woman | 8 | 27.6% | boat | 4 | 26.7% | landscape | 6 | 35.3% |
| 3 | sharks | 11 | 34.4% | table | 5 | 17.2% | American | 4 | 26.7% | peasants | 3 | 17.6% |
| 4 | sea | 9 | 28.1% | portrait | 5 | 17.2% | water | 3 | 20.0% | grain | 3 | 17.6% |
| 5 | ocean | 9 | 28.1% | gown | 4 | 13.8% | ice | 3 | 20.0% | wheat | 3 | 11.8% |

Chart 5: steve tagger: consistency in terms across known users; Top five tags for each of the top four works

4.7 Terms assigned to the other works in the top four

The tags assigned to the three other ‘most tagged’ works in the steve tagger prototype (Chart 6, Chart 7, and Chart 8) share many of these characteristics.

17th SIG/CR Classification Research Workshop, November 4, 2006

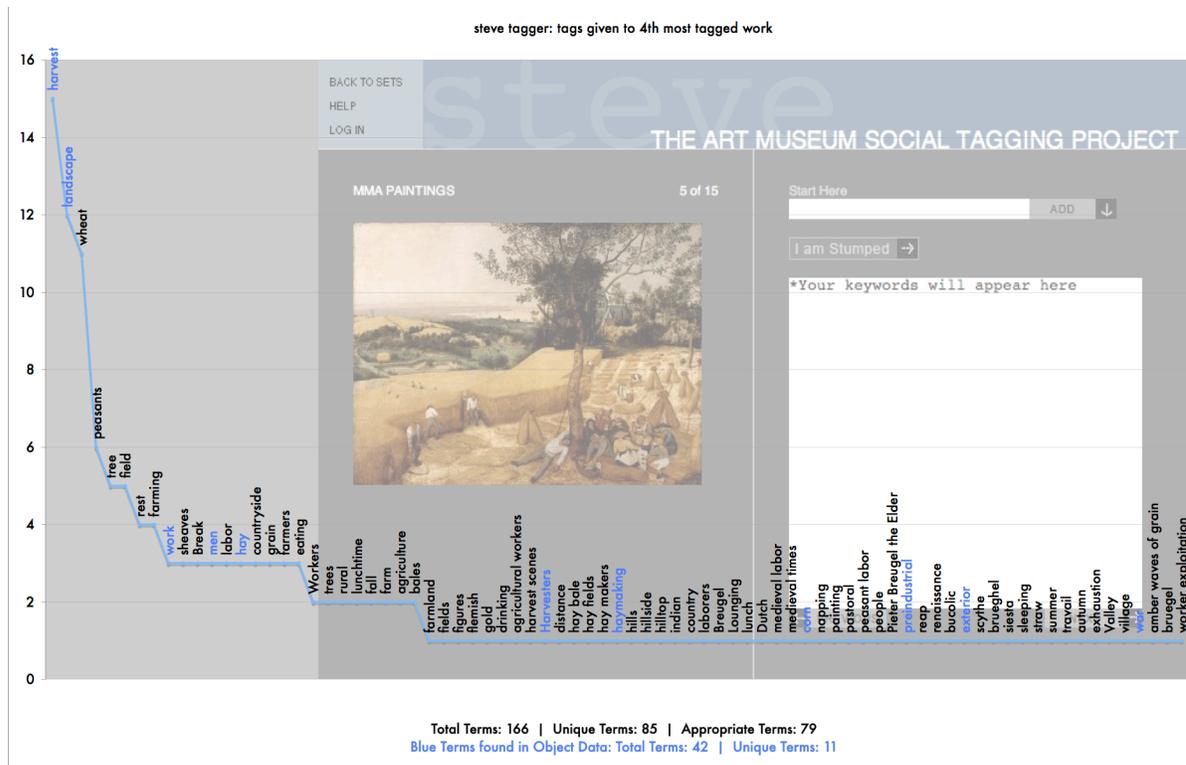


Chart 8: Tags assigned to the fourth most tagged work in the steve tagger: Pieter Bruegel the Elder (Netherlandish, active by 1551, died 1569), *The Harvesters*, 1565, Oil on wood; Overall, including added strips at top, bottom, and right, 46 7/8 x 63 3/4 in. (119 x 162 cm); original painted surface 45 7/8 x 62 7/8 in. (116.5 x 159.5 cm), *The Metropolitan Museum of Art, New York, Rogers Fund, 1919* (19.164)

A small number of terms are assigned frequently, tapering out into a large number of terms assigned only once. Agreement on the first group of terms is high (Chart 5): 47.1%-31.0% of known taggers assigned the most common terms to these works. But for two of the works, there is a sharper divergence: the fifth-most common term was assigned less than 15% of the time for



For all works, a significant portion of the terms assigned in the steve tagger prototype were not found in the museum's on-line documentation for these works: an average of 90.2% of the terms were new. (Chart 3 shows these details for all works.) There are differences among the works in whether the new terms were highly occurring ones or not, pointing to a need to study the

variations in tags assigned to different genres of works of art. For example, the most common term for the history painting “washington” was a term found in the documentation, for the title of this work is *George Washington Crossing the Delaware*. However the title of *The Gulf Stream*, a more allegorical work, does not refer directly to its subject matter.

5. Professional and Public Vocabularies

The comparison of tags assigned to the top four most-tagged works in the steve tagger prototype with their documentation on the Web site of The Metropolitan Museum of Art, confirms the distinction between public and professional vocabularies pointed to by the Proof of Concept studies (Trant, 2006). A review of the curatorial notes for these works, confirms that, to paraphrase a curator at The Metropolitan Museum of Art, everything the specialist knows is not in the picture (Jenkins, 2006). These texts discuss the price of the paintings (2), their critical reception (2), the process of their creation and existence of studies (2), and their significance in the history of Western painting (1). In art historical discourse subject matter, unless it is problematic, is often taken as given; the work is assumed to be present and visible. (These genres are explored further in the templates created for the Pachyderm project (Samis & Johnson, 2005).) When we change the role of texts from their original purpose of supporting interpretation to the additional support of information retrieval, they may not be as effective. Social tagging seems a promising way to supplement museum records with terminology to support some kinds of queries.

6. Research Questions

Many questions remain about the effectiveness of social tagging and folksonomy to enable access to art museum collections (Figure 8). Within the steve collaboration we’re committed to building our understanding of the way people tag, and the nature of the resulting folksonomy. For example, a comparison of tags to lexicographical resources like wordNet could identify clusters of like terms in the tags and speed our analysis of the use of synonyms, different parts of speech and descriptions of subject matter at different levels of specificity. Comparisons with disciplinary resources like the Art and Architecture Thesaurus and Union List of Artists Names will show if tags represent concepts from the art historical domain, or if they are introducing new ways of looking at art. Comparisons of terms with the searches (successful and unsuccessful) of museum databases will help us establish whether the folksonomy that results from social tagging could

fulfill a need in on-line retrieval of art museum collections. We also need to develop effective ways of determining if tags are appropriate to works of art, so that we can effectively process user-supplied tags. These questions are now being addressed in the context of a larger-scale research project headquartered at The Metropolitan Museum of Art and funded by the Institute of Museum and Library Services (IMLS).

social tagging and folksonomy analysis

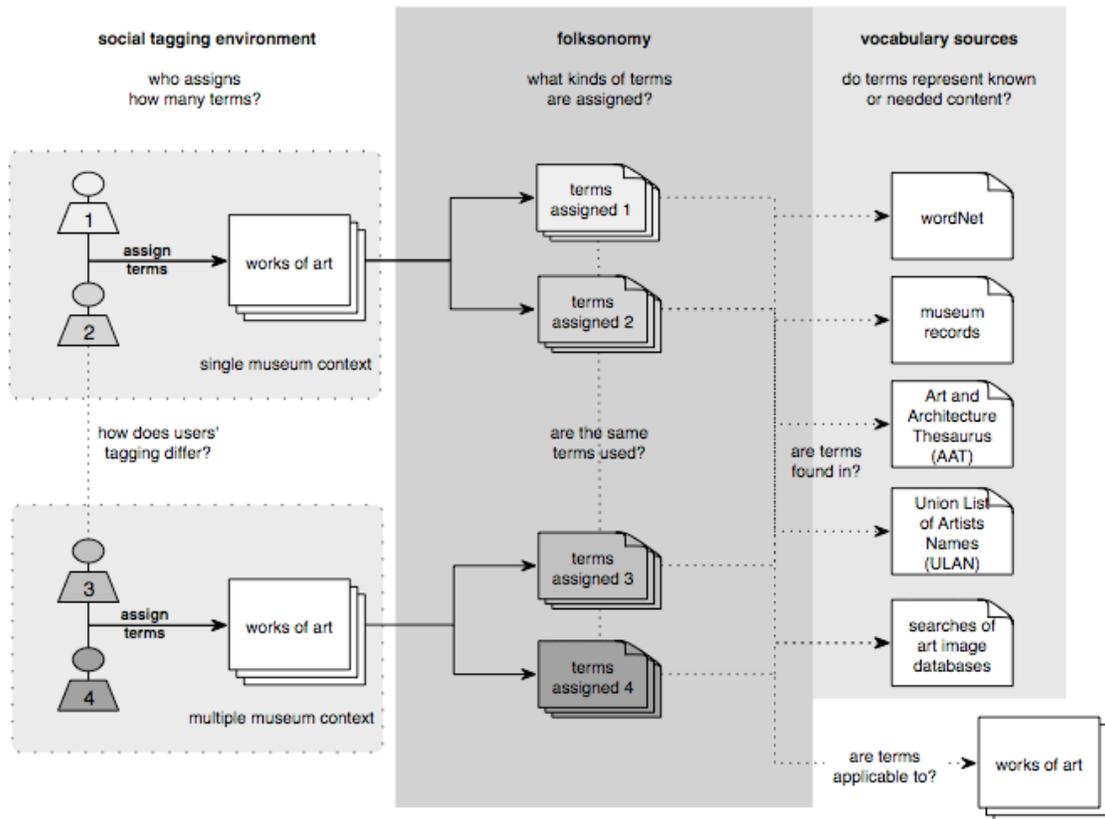


Figure 8: Research questions are positioned in the data collection and analysis process, from term collection in a social tagging environment, through folksonomy analysis, comparison with controlled vocabularies, and assessment in relationship to the work of art

7. Affordances of Social Classification

The results of the Proof of Concept and Prototype studies give us cause for optimism: social tagging and folksonomy could improve the accessibility of museum collections by enhancing information retrieval. But there may also be some additional benefits of these systems in the museum environment. Tagging is a highly personal activity (Golder & Huberman, 2005). Tags

exist in a liminal space between a user and an information resource, and as such represent a critical facet of personal meaning-making. The subjective nature of tagging might reveal something of how art collections are perceived by a broad public. There are few tools to gather this kind of feedback directly from museum visitors; the social side of social tagging merits exploration within the context of other museum-based community development and user-contributed content initiatives. Social tagging offers new way for museums to engage user communities and, through the resulting folksonomy, to assist them in their use of collections.

The use of literary warrant evolved as a way to verify terminology used in controlled vocabularies against language (and concepts) in use in a community during the construction of the *Art and Architecture Thesaurus*. Social tagging offers a view of the language of a different, more broadly public user group, one that art museums also wish to engage. The diversity represented in tags may help the museum meet the needs of the many communities interested in – and represented in – art museum collections. Integration of folksonomic strategies into on-line art museum documentation seems a promising way to supplement scholarly and professional perspectives, and support the multiple points of view represented by on-line users. It certainly merits further, serious study, both of the terminology resulting from social tagging (as discussed here) and of the social nature of tagging systems. A measured approach has been taken within steve, to prove the value of public contributions prior to deploying systems linked to museum on-line catalogs. Others, such as the Powerhouse Museum (Powerhouse Museum, 2006; Powerhouse Museum & Chan, 2005) and Smithsonian Photography (Smithsonian Institution, 2006) are moving ahead in a single-institution context, deploying systems that engage users in tagging collections, and enabling searching based on the resulting folksonomy. These initiatives, and others that engage museum visitors on-line and on-site will add to our understanding of the complex relationship between museum and museum-goer, between institution and individual, and between information and personal interest.

References

Ahn, L. v., and Dabbish, L. (2004). Labeling images with a computer game. In *CHI 2004: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (pp. 319-326). Vienna, Austria: ACM.

Barnard, K., Duygulu, P., and Forsyth, D. (2001). Clustering art. In *CVPR 2001: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (vol. 2, pp. 434-439).

Bearman, D., and Trant, J. (1997). *Museums and the Web 97: Proceedings*. Pittsburgh, PA: Archives & Museum Informatics.

Bearman, D., and Trant, J. (2005). Social terminology enhancement through vernacular engagement: Exploring collaborative annotation to encourage interaction with museum collections. *D-Lib Magazine* 11.

Benvenuto Di Giovanni, I. S., 1436-c.1518 (Artist). (1470s). *Expulsion from Paradise*.

Canadian Heritage Information Network (CHIN). (2002, 2003-02-10). *CHIN Data Dictionary: Humanities*. Retrieved November 9, 2005, from <http://daryl.chin.gc.ca:8000/BASIS/chindd/user/wwwhe/SF>.

Cataloging by Crowd Working Group, and Leonhardt, T. (2005, July). Community cataloguing application: Description of V1.0 prototype. Retrieved August 24, 2005, from <http://www.steve.museum/reference/CCA-PrototypeLeonhardt200507.pdf>.

Cataloging by Crowd Working Group, and Trant, J. (2005, July 23). Cataloging by crowd: Issues in term analysis. Retrieved August 25, 2005, from <http://www.steve.museum/reference/comCatTermAnalysis050723.pdf>.

Chun, S., Cherry, R., Hiwiller, D., Trant, J., and Wyman, B. (2006). Steve.museum: An ongoing experiment in social tagging, folksonomy, and museums. In J. Trant and D. Bearman (eds.),

Museums and the Web 2006: Selected papers from an international conference. Albuquerque, NM: Archives & Museum Informatics.

Cunningham, S. J., and Masoodian, M. (2006). Looking for a picture: An analysis of everyday image information searching. In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries* (pp. 198-199). Chapel Hill, NC: ACM.

Doolan, J., Peacock, D., and Ellis, D. (2004). Searching for meaning, not just records. In D. Bearman and J. Trant (eds.), *Museums and the Web 2004: Proceedings*. Washington, DC / Arlington, VA: Archives & Museum Informatics.

Elinich, K. (2004). One Wright Way: From collections to classrooms. In D. Bearman and J. Trant (eds.), *Museums and the Web 2004: Proceedings*. Washington, DC / Arlington, VA: Archives & Museum Informatics.

Golder, S., and Huberman, B.A. (2005). Usage patterns of collaborative tagging systems. *Journal of Information Science* 32: 198-208.

Hammond, T., Hannay, T., Lund, B., and Scott, J. (2005). Social bookmarking tools (I): A general review. *D-Lib Magazine* 11(4).

Harpring, P. (2002). The language of images: Enhancing access to images by applying metadata schemas and structured vocabularies. In M. Baca (ed.), *Introduction to art image access: Issues, tools, standards and strategies*. Los Angeles: J. Paul Getty Trust.

Honigsbaum, M. (2005, December 12). Search for gay history to create "virtual museum". *The Guardian*. Retrieved October 10, 2006, from <http://education.guardian.co.uk/higher/research/story/0,9865,1665298,00.html>.

Janney, K., and Sledge, J. (1995, September). Formulating typical questions. In *A User Model for CIMI Z39.50 Application Profile*. Retrieved May 3, 2006, from http://www.cimi.org/old_site/documents/Z3950_app_profile_0995.html.

Jenkins, M. (2006, January 10). Everything I know isn't in the picture. Personal communication.

Kydd, S., Mackenzie, D., and Myles, M. (1998). Multiple uses of data in the museum environment. In D. Bearman and J. Trant (eds.), *Museums and the Web 1998: Proceedings*. Toronto, ON: Archives & Museum Informatics.

Lund, B., Hammond, T., Flack, M., and Hannay, T. (2005). Social bookmarking tools (II): A case study: Connotea. *D-Lib Magazine* 11(4).

Markey, K. (1984). Interindexer consistency tests: A literature review and report of a test of consistency in indexing visual materials. *Library & Information Science Research* 6: 155-177.

Marlow, C., Naaman, M., Boyd, D., and Davis, M. (2006). HT06, tagging paper, taxonomy, Flickr, academic article, to read. In *Proceedings of the 17th Conference on Hypertext and Hypermedia* (pp. 31-40). Odense, Denmark.

Mathes, A. (2004, December). Folksonomies: Cooperative classification and communication through shared metadata. Retrieved August 6, 2005, from <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>.

McCorry, H., and Morrison, I. O. (1993). *Report on the Catechism Project*. Edinburgh: National Museums of Scotland.

Pennsylvania State University Library. (2005-). PennTags. Retrieved April 15, 2006, from <http://tags.library.upenn.edu/>.

Pink, D. (2005, December 11). Year in ideas: Folksonomy. *The New York Times Magazine*.

Powerhouse Museum. (2006). OPAC 2.0. Retrieved October 10, 2006, from <http://www.powerhousemuseum.com/collection/database/>.

Powerhouse Museum, and Chan, S. (2005, ongoing). Electronic swatchbook. Retrieved January 10, 2006, from <http://www.powerhousemuseum.com/electronicswatchbook/>.

Quintarelli, E. (2005, June 24). Folksonomies: Power to the people. Paper presented at the ISKO Italy / UniMIB Meeting, Milan.

R. W. De F., and The Metropolitan Museum of Art. (1910). Expert guidance to the museum. *Bulletin of The Metropolitan Museum of Art* 5: 201-207.

Reich, C., and Lindgren-Streicher, A. (2006). Visitor interactions with digitized artifacts. In D. Bearman and J. Trant (eds.), *Museums and the Web 2006: Proceedings*. Albuquerque, NM: Archives & Museum Informatics.

Reilly, B. (ed.). (2000). *Collections: Museum collections online*. Council on Library and Information Resources.

Roberts, H. E. (1994). Second hand images: The role of surrogates in artistic and cultural exchange. *Visual Resources* 9: 335-346.

Samis, P., and Johnson, L. (2005). Taking teaching by the tusks: Introducing Pachyderm 2.0. In D. Bearman and J. Trant (eds.), *Museums and the Web 2005: Proceedings*. Vancouver, BC: Archives & Museum Informatics.

Samis, P., Mitroff, D., and Johnson, T. (2005). Of Ansel and Atomz: Surfacing deep content on-line and on-site at SFMOMA. In D. Bearman and J. Trant (eds.), *Museums and the Web 2005: Proceedings*. Vancouver, BC: Archives & Museum Informatics.

Schaller, D., and Allison-Bunnell, S. (2003). Practicing what we teach: How learning theory can guide development of online educational activities. In J. Trant and D. Bearman (eds.), *Museums and the Web 2003: Proceedings*. Charlotte, NC: Archives & Museum Informatics.

Shubert, S. B. (1996). Subject access to museum objects: Applying the principles of the subject approach to information from library and information science to the documentation of humanities museum collections. Ph.D. dissertation. University of Toronto.

Sledge, J. (1995). Points of view. In D. Bearman (ed.), *Multimedia computing and museums: Selected papers from the Third International Conference on Hypermedia and Interactivity in Museums* (ICHIM95/MCN95; pp. 335-346). Pittsburgh: Archives & Museum Informatics.

Sledge, J., and Case, M. (1995). Looking for Mr. Rococo: Getty Art History Information Program Point of View Workshop. *Archives and Museum Informatics* 9: 124-129.

Smithsonian Institution. (2006). Smithsonian photography. Retrieved October 10, 2006, from <http://photography.si.edu/>.

Stephenson, C., and Mcclung, P. (eds.). (1998). *Delivering digital images: Cultural heritage resources for education*. Los Angeles, CA: Getty Research Institute.

Trant, J. (2006). Exploring the potential for social tagging and folksonomy in art museums: proof of concept. *New Review of Hypermedia and Multimedia* 12: 83-105.

Trant, J., and Bearman, D. (eds.). (2006). *Museums and the Web 2006: Selected papers from an international conference*. Toronto: Archives & Museum Informatics.

Trant, J., and Wyman, B. (2006). Investigating social tagging and folksonomy in art museums with steve.museum, In *World Wide Web 2006: Tagging Workshop*. Edinburgh, Scotland: ACM.

Vander Wal, T. (2005, November 2). Folksonomy definition and Wikipedia. Retrieved October 10, 2006, from <http://www.vanderwal.net/random/entrysel.php?blog=1750>.

Vemuri, N. S., Torres, R. D. S., Shen, R., Gonçalves, M. A., Fan, W., and Fox, E. A. (2006). A content-based image retrieval service for archaeology collections. In *ECDL 2006: Proceedings of the 10th European Conference on Research and Advanced Technology for Digital Libraries* (Lecture Notes in Computer Science, no. 4172; pp. 438-440). Berlin: Springer.

Wang, J. Z., Li, J., and Lin, S. C. (2003). Evaluation strategies for automatic linguistic indexing of pictures. In *IEEE International Conference on Image Processing (ICIP)*. Barcelona, Spain.

Wang, J. Z., Li, J., and Wiederhold, G. (2001). SIMPLIcity: Semantics-sensitive integrated matching for picture libraries. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23: 947-963.

Ward, A., Eliot, N., Graham, M., Riley, J., Sheen, N., Eakins, J., et al. (2001). Collage and content-based image retrieval: Collaboration for enhanced services for the London Guildhall Library. In J. Trant and D. Bearman (eds.), *Museums and the Web 2001: Proceedings*. Seattle, WA: Archives & Museum Informatics.