Browsing Images Using Broad Classification Categories

C. Olivia Frost
Anna Noakes
School of Information
University of Michigan
Ann Arbor, Michigan

Abstract

Browsing can serve an important role as a strategy for retrieving image information and can allow users to recognize what they are looking for more easily than describing it. This article describes a project which developed and tested an art history image data base which used broad classification units to group image sets into meaningful categories that support browsing. Categories from the Art and Architecture Thesaurus were used to provide broad classification units to guide users in browsing and searching. The system provided both browse and search capabilities and was tested with art history specialists and generalist users. The type of user and the use to which the image is put will influence the type of search strategy employed. Preferences for, and performance in, the browse and search modes varied according to the user's familiarity with the subject domain. Both sets of users felt that the ability to use both browse and search in a multimodal system would best meet their search needs.

Introduction

We describe a research project which utilized classification as a means to group image sets into meaningful categories that support browsing. A key premise in the research is that browsing can serve an important role as a strategy for retrieving image information. As Kwasnik has observed, browsing relieves the user from the burden of formulating a precise search strategy and capitalizes on the fact that it is easier to recognize what is interesting or useful than it is to specify it in advance. A related advantage is the ability to facilitate exploration without prior knowledge of subject content, and thus enable users to cross over into unfamiliar domains (Kwasnik 1992). Furthermore, browsing lends itself particularly well to visual images since a pictorial image is able to present itself whole (Arnheim 1970) and in its own medium of expression.

With this premise in mind, researchers at the University of Michigan School of Information (SI) developed and tested a text-based image retrieval system called the SI Art Image Browser (http://www.si.umich.edu/Art_History/) which used broad classification units to group image sets into meaningful categories that support browsing. While textual descriptors have their limitations, they can be used in a way that allows users to employ their cognitive abilities to scan image content within sets of images to retrieve desired information.

The browsing approach employed in the SI Art Image Browser provides the user with access to image sets in addition to individual images. Shatford Layne (1994) points out advantages to this type of approach - a user searching for a single image may be able to make a selection by browsing through a group of images. Abstract or subjective concepts are difficult to convey in a textual description. Or, scanning a selection of images may prove more efficient than indexing at a high level of detail. Shatford Layne further argues that "the emphasis when indexing images should be on recall, rather than on precision. Rather than devoting time to extraordinarily detailed or complicated indexing ... it might be better to concentrate on indexing the basic elements of an image and rely on scanning, or browsing, to make the fine distinctions." (p. 586) For our system, we relied on broad classification categories rather than a highly detailed scheme.

Our study sought to determine the usefulness of classification as a retrieval device for images from a digital database in a networked environment. The research team designed, implemented and evaluated an experimental image browser system on the World Wide Web which uses classification to facilitate browsing of art images. A database of over 2,000 digitized images was built from three collections at the University of Michigan and a private collection. Once the system was designed and implemented, it was evaluated in a laboratory setting and used people with real information needs which could be satisfied by the images.

Our image retrieval system offers users two approaches, which can be used in a complementary fashion: 1) structured browsing using classification structures as an intellectual framework to help navigate the collection; and 2) direct term search using a keyword term. In both approaches, a successful search will result in a set of images which can be browsed. The premise underlying both approaches is that with images, it is often easier to recognize what is wanted than to try to formulate a query, and that one of the most effective ways to find images is by browsing a set of image surrogates which are representations of the image itself. For these reasons, thumbnail size images can provide a useful surrogate of the full size image and give the user enough information to make a judgment as to whether a given image will be useful or not. In the approach using a classification framework, the user can be guided from broad to progressively narrower subtopics, and can progress through the category tree until he or she gets to a point where the sets of images are located.

The research also emphasizes the user perspective in the design and testing of the system, and, in developing the design of the browser, the project team used an iterative approach which questioned art history, art and architecture faculty, librarians and students to determine patterns of use and information searching practices.

Progress in image retrieval
Images can be indexed and retrieved by textual descriptors and by image content. In textual queries, words are used to retrieve images or image sets, and in visual queries (content-based retrieval), images or image sets are retrieved by visual characteristics such as color, texture, and shape. Image retrieval strategies based on either text-based or content-based retrieval alone have their limitations. Images are often subject to a wide range of interpretations, and textual descriptions can only begin to capture the richness and complexity of the semantic content of a visual image. Human indexing of images is highly labor-intensive and limiting when large databases are involved. However, retrieval based on visual characteristics is computationally

intense and has not yet reached the point where it can be efficiently used to formulate intellectually subtle queries, especially for non-specialist users.

Research in content-based image retrieval includes query-by-visual sample, pictorial queries, and specification of image features. Many systems use a combination of types of image queries, and others combine textual information with image contents to retrieve images. The IBM QBIC system provides both visual sample-based queries and sketch-based queries. Holt and Hartwick (1994) report on the usefulness and effectiveness of searching art images in a system which classifies visual elements in an image using the QBIC software.

User-oriented studies

A number of studies have explored what goes into a user's query. Hastings' dissertation research (1994) looked at categories for types of information needs of art historians as articulated in discussions by Markey (1988) and Stam (1989), and then analyzed queries by art historians searching a database of digitized art images to determine changes that occur in the query when the surrogate image format changes from photograph to digital. In a later study (1997), Hastings extended the research to a distributed network of digital art images, and used survey data, interviews, and query analysis to examine categories of queries and types of index access points used in searches. A study by Ester (1991) also focused on art historians and their viewing of surrogate images of art objects, with a view towards determining thresholds between higher and lower quality images in terms of resolution and range of colors. Enser (1995) conducted a study of user-assigned image descriptors from a survey of requests for still image material received by a major picture library. Analysis of this request collection addressed the categorization of requests by subject and user type.

Browsing

A number of researchers have advocated the inclusion of browsing features into the design of information systems including Kwasnik (1992) and Bates (1989). Bawden (1993) and Chang (1993) provide useful overviews of the role of browsing. Few studies have focused on browsing as a strategy for retrieval of images. Besser's prototype system used visual browsing tools on computer workstations to effect retrieval of visual images (Besser 1990, 1992). From the database research perspective, Grosky stresses the importance of the browsing paradigm in managing multimedia data (Grosky 1994), and develops a strategy to allow users to utilize metadata to intelligently browse through a collection of media objects (Grosky et al. 1994).

Research design

The SI Art Image Browser project work centered on activities needed to: 1) build a database of images; 2) design, build, and implement a system; 3) identify classification categories; and 4) evaluate the resulting system with real users using an iterative approach. The Art Image Browser was evaluated using four approaches:

1) **Ease of use study**. An ease-of-use study using domain and non-domain experts was conducted prior to the initial release of the browser. Its primary purpose was to test users'

ability to navigate successfully through the browser; results collected from the study led to further refinements to the interface before additional studies were conducted.

2) **Quasi-controlled study**. Once the final version of the browser had been released, a quasi-controlled study was conducted to determine whether statistically significant differences existed in the comparative "success" of subject domain experts and nonexperts when using the browser for image retrieval. An important qualitative component to the study tried to further determine a predilection for searching or browsing within such a system and to explicate the differences between the searching behavior of experts and nonexperts.

3) **Focus group interviews**. A series of focus group interviews comparing the responses of subject domain experts and nonexperts provided a rich source of qualitative data in areas ranging from general uses for art images to information seeking; text based information users would like to see accompanying online images; and specific features subjects would most like to see incorporated into an image delivery system.

4) **Online survey**. An online survey was conducted in order to gather feedback on the browser from users of the World Wide Web and to seek additional information about user search behavior. Information gathered confirmed many of the findings in the other studies and confirmed the success of many of the browser's technical capabilities.

This paper primarily discusses the classification design and the quasi-controlled study.

Research questions

The research was focused on the use of broad classification categories to provide structured browsing for sets of images. We were interested as well in the browse and search modalities in image information seeking, and the implications for classification as a means to navigate a search space of images. In designing and testing a browsing prototype we focused on questions which included:

> *Is a browsing model, in which images are classified into navigable hierarchies, a useful strategy for providing intellectual access to digitized images? How well would users be able to locate specific images or groups of images using this system? Are there performance differences between the three search modes of browse, search, and a combination of the two?*

> *To what extent can we rely on visual image information alone in browsing images? What types of accompanying textual information do users need in their search for images?*

The type of user and the use to which the image is put will influence the type of search strategy employed. In the design of our prototype systems, we considered the idea of separate systems for expert and generalist users. The version of the browser which became our eventual prototype was designed to facilitate image searching by domain specialists as well as generalist users. As part of our research, we conducted a quasi-controlled test to establish the presence of statistically significant differences between art domain experts and nonexperts when browsing, searching, and

locating art images using the prototype SI Art Image Browser. The test collected both quantitative and qualitative data to answer questions regarding differences in speed, accuracy, strategy, cumulative performance and system preference that may exist between the two groups of participants. We asked:

> *Would a structured browsing approach be useful for a broad array of users interested in art images - both generalist users as well as subject domain specialists? Would a classification approach be useful to expert users, who would be more likely to have the domain expertise to search directly?*
>
> *Can a single set of classification categories serve the needs of art history experts as well as generalist users?*

User-centered design and an iterative approach

Since our system was developed with the needs of the end-user in mind, feedback from present and potential users played a key role in both the development of the classification design as well as the system design and interface. The project team used an iterative approach and questioned art history faculty and students - as well as librarians - to determine patterns of use and information searching practices.

To assist in the design of the browser, we first created two pilot systems to solicit feedback from users and used this input to inform later versions of the system. The first pilot system was demonstrated to a group of about 100 students, faculty, and staff from the School of Information and other units of the University, and informal feedback was solicited from the viewers.

In addition, interviews were held with faculty from the History of Art Department, the School of Art, and the College of Architecture and Urban Planning, as well as focus group sessions held with students in an art history class. The interviews were informal discussions which centered on the demonstration model, and which elicited helpful information about users' impressions of the database structure, classification categories, interface design, and image quality. This feedback was used to inform later versions of the prototype.

Classification design

In developing classification structures for browsing, the project drew from established models used in the subject disciplines of art and architecture. The project team also worked with domain specialists in art and architecture to investigate the suitability of models including the Library of Congress Classification, the Dewey Decimal Classification, the Fogg classification schemes, a modified Metropolitan Museum of Art classification system, and the Art & Architecture Thesaurus. A wide variety of classification schemes have been used for the organization of art materials, and the majority of these classification systems have a similar basic outline, which arranges works according to facets such as medium/technique, style, period, geographic or cultural area, artist, type of building, subject, date, and title. Based upon information drawn from the interviews with users about their information needs and information seeking practices, and working from the subject classification models drawn from classification examples, the project team tried out design prototypes and solicited feedback from users. The team was assisted by a group of six students from the School of Information in a class in Organization of Nontextual

Information Resources; students in this class used the research as a vehicle to explore classification and metadata for image information. For their term project, the students developed a working model and solicited evaluations from prospective users.

Our initial plan was to provide one system for domain (art history) specialists and another for general users. Our first prototype database system allowed us to test whether this dual approach is appropriate, or whether one system can be modified to suit both sets of users. Responses from non-domain and domain users led us to conclude that a single system could suit both audiences. In discussions with students and faculty in art history and fine arts, we were told that the idea of using a broadly conceived browsing tool to create subsets of the image database, with the capability to create thumbnail galleries based on fairly specific criteria, would be an effective way to encourage use of the database by art history professionals.

The first browser pilot contained category areas for country, general medium, specific medium, decade, specific year, technique, artist name, title of work, and content or identifiable object. We decided to test the feasibility of using the Dewey Decimal Classification (DDC) scheme to guide browsers by subject through a collection of digital images. Student team members set to work assigning DDC subject categories to 100 images, and were attracted to its hierarchical structure and flexibility. However, most of the images in this database were 20th-century art images, and the students found DDC to be inadequate in its ability to handle the description of abstract images.

A second pilot was developed which incorporated interface improvements suggested by the first group of faculty consultants. The content area of the first pilot was 20th-century art. The selection of this content was determined largely by the immediate availability of this database; however, the database proved unsuitable due to the poor quality of its images, and we turned to a second source of images which became available - a set of slides containing over 200 images of Italian architecture.

In developing the second pilot, we examined the art documentation and classification scheme employed by the Canadian Centre for Architecture, with which one of our team members had had prior experience. This classification scheme provided further input from which we developed categories and browsing hierarchies for the prototype of architecture images. We decided on the categories of Location, Title, Chronological Period, Stylistic Period, Architect, and Purpose/Building Type. Since our eventual database was broader than architecture, we later adapted some of these categories for more general use.

Collection development constraints

Collection development issues played a part in shaping our classification strategies, and we faced numerous challenges in the development of a collection. Since our classification scheme was to be used in a database of Internet-accessible images, copyright constraints affected the degree to which we could provide the desired content depth and balance of our image collections. We were using metadata provided by our museum partners, and this metadata had been developed in-house over a number of years and without external guidelines for standardization of metadata elements. The information had been largely for internal use and was certainly not designed for the purpose of guiding generalist users in browsing surrogate images online. Thus, metadata-data

control issues often slowed the development of the database and limited the categories which could be used in building the data fields. We found that we had to work within the bounds of categories contained within the museum curators' own cataloging data.

Ultimately, the determination of the classification categories and structure was most heavily influenced by the characteristics of the data provided by the various content providers. Ideally, our prototype classification design would be broad-based, and our collections would be broad in scope and even in the distribution of images populating the various categories. We sought to achieve content balance and depth, to appeal to a broad set of users, and to encourage comparison of images. That was our ideal - practically speaking, however, the project's collection content in its early stages was in large part driven by constraints imposed by copyright issues and the availability of suitable metadata. Our initial prototypes were thus narrower in scope than their eventual successor, which evolved at a point at which the copyright and metadata issues had been resolved.

Classification scheme used for testing the browser

The *Art and Architecture Thesaurus* (1994) served as our major source for broad categories to be used in the eventual system used for formal testing of the browser. It also served as a good source for the standardization of descriptive terminology and for notes to assist non-domain users in understanding some discipline-specific terminology. Working from the *Art and Architecture Thesaurus* (AAT), we selected seven facets for use in the browser. These broad conceptual categories serve as a starting point for users to begin their browsing. The categories also serve as fields within which the user can execute a search for a particular artist, title, etc. By combining terms at the point of the search, the broad categories can be assembled by the user to formulate more precise search terms to be used for more specific browsing or a specific known-item search.

We used the AAT category of *Agents* for artist's names. The *Space* category in our system groups images according to the nationality of the artist. We provided a world map to enable users to click first on the continent of interest, and then proceed to a breakdown of time categories within the geographic location. *Time* is used to break the content into chronological sub-divisions which vary in their level of specificity depending on the depth of material in a time segment. *Materials* was used to designate medium; i.e., what the material is made of or the process used to make it. *Objects* was used for type of artistic creation, e.g., painting, sculpture. The categories under *Objects* and *Materials* usually generated large sets which we wanted to further subdivide topically. Team members with backgrounds in history of art were assigned the task of determining subcategories to break out the large sets into more manageable categories. However, this effort proved too time-consuming to accomplish within the standards the team members had set for themselves and the deadline of the system production.

We added a *Subject* category to describe objects depicted within the image, but not necessarily the theme of the image. Subject or concept categorization poses a particular challenge for art, where intrinsically there are different levels of subject meaning. Panofsky's (1962) categorization of three levels of subject matter into primary, secondary, and iconological has served as a useful basis for distinguishing among levels of meaning. Since the content providers had not followed a standardized list of terms in the assignment of subject terms, the vocabulary could not be controlled. The lack of universally-accepted standards and vocabulary control for art

documentation remains a barrier to sharing data about artworks and cultural materials between institutions. One group within the project team sorted out the list of subject terms from the original list to arrive at a small set of frequently occurring broad categories considered to be of general interest to users; e.g., *Animals, People, Places*. This set of categories was determined by the non-expert group members, who probably felt less constrained than the domain members of the team in arriving at a general set of topical categories.

Styles and periods is an AAT category type that we had used as a category in the second browser prototype and hoped to include in the final version as well. However, the museum curator content providers felt that these designations could not be assigned with confidence and accuracy, since many works had overlapping or ambiguous periods. The art history community appears to disagree on the exact date ranges for these various stylistic periods, and individual curators may be reluctant to assign a given work to a stylistic category. Curators also argued that many works of the same chronological period are not readily classified this way, even when created by the same artist. Further arguments were that information about stylistic period was not recorded in the image providers' databases; and that stylistic information is already common knowledge among art domain experts. While our project team still remained convinced that stylistic period categories, limitations notwithstanding, would be a valuable browsing component for general users especially, we recognized that this would not be feasible, given the opposition by the curators providing the image content and the metadata. As had occurred in the issue of breaking down the *Objects* and *Materials* categories into manageable sub-units, there was a strong difference of opinion between the subject experts, who were concerned about how potentially controversial decisions on categorization would be viewed, and the general project goals of providing manageable browsing guides to assist generalist users.

Browser System Design
The system we designed has both a browse and search modality. In choosing the browse approach, the user begins by making a selection from a menu organized by broad conceptual categories (artist, title, date, medium, object type, and subject), which takes the user to listings of subdivision groupings which in turn lead to sets of thumbnail images. The user can then click on the thumbnail to retrieve a document containing the thumbnail and textual information pertaining to the image. This document page also leads to the full-page image.

At any point in the browse stage, users can choose a "search" option and execute a query in any of the category fields of artist, medium, date, etc. Users can also search on a combination of these categories. These queries can be used to generate image sets which can in turn be scanned. The project developed a search engine for searching textual descriptions and access points. The FTL retrieval application (Full Text Lexicographer) provides enhanced search capabilities and supports several kinds of searches: words and phrases, Boolean operations, word and sentence proximity, and stemming. FTL allows searches on any of the fields in the database, and the user is able to ask, for example, what the collection has by a given artist within a specific time span using a specific medium.

The browsing categories of *Artist, Nationality, Title, Object Type, Medium, Date* and *Subject* constitute the top level categories in the Browser. Each category at the first level is associated with a number of sub-categories, which comprise a second level. Each of these subcategories is

further subdivided into terms which form the nodes of a third level. Each of these nodes is associated with a set of images. If the node is chosen, the thumbnails of the first ten images are displayed. At this point, the user can choose either to see more thumbnail images or select a particular thumbnail image. When a thumbnail image is selected, the description of the image will be displayed, from which a user can also select the full size image. In the example below, if a user chooses the subcategory *Places* from the *Subject* category, then selects *Venice*, the thumbnail images shown in Figure 1 will be displayed on the screen.
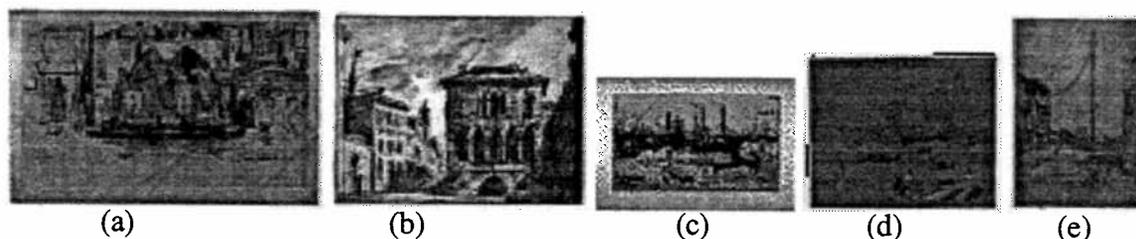

(a)   (b)   (c)   (d)   (e)

Figure 1: Thumbnails of four images of Venice.

If a user clicks on the image (b) in Figure 1, she will see the display of the description of the image as shown in Figure 2.

Street in Venice

Location: Ann Arbor, MI
Artist: Anonymous, Italian
Nationality: Italian, Venetian
Artist Dates:
Title: Street in Venice
Date: 18th century
Object Type: drawing
Medium: brown ink, gray wash and black chalk
Subjects: Venice, buildings, cityscape, soldiers
Holding Institution: The University of Michigan Museum of Art
Accession No.: 1970/2.14
Credit Line: Gift through the Estate of Edward Sonnenschein
Image Source: The University of Michigan Museum of Art
Image Accession No.: PCD4097-2012-0833-21

Figure 2: Selected image and description.

If the user clicks on the thumbnail image on this display, the full image will be displayed.

In the more direct search mode, users can enter a keyword or words in any of the seven categories used for the browsing structure (*Artist, Nationality*, etc.)   Once a match is found, the images associated with the particular category are displayed. If more than one keywords are given, the system performs a logic operation "AND" on the images, and returns with the images whose description matches all the keywords given by the user. For example, if a user gives the keyword of *Venice*, she would get the images shown in Figure 1. If the user types in *Venice* and *street* , the system will return a thumbnail image that links to the detailed information shown in Figure 2. As in the browsing example, the user can retrieve the full image by clicking on the thumbnail image.

Evaluation design

The research design for evaluation of the SI Art Image Browser was intended to elicit differences in the use of the browser between domain experts and nonexperts, as well as differences among the two groups in their use of the three distinct search modes: browsing, direct search and the hybrid mode combining browsing and searching modalities.

Prior to task and system assignment the research team grouped subjects into two classes by their level of art-domain expertise. The expert group was comprised of members from the local university and museum environments; specifically, faculty, lecturers, and advanced graduate students in art and design, art history, and architecture and urban planning. We also solicited the participation of curatorial staff in art and archaeology museums and in the art history slide and photograph collection. Non-expert subjects were recruited from various university residences and academic departments.

The research team members scheduled subjects for the controlled study at a variety of times during the day over a period of two months. Test ádministrators met the subjects at mutually agreed-upon locations which included campus computing sites, campus libraries, university residences, offices, and private homes. The greatest number of distributed locations was selected for test purposes as the prototype system was intended for distributed network computing. However, a minimum modem speed of 24.4 baud was required in order to minimize the impact of system speed on the subjects' time to complete tasks.

We provided subjects with a general description of the project and ascertained their basic familiarity with computers and the Netscape Web browser. If subjects were unfamiliar with Netscape, they were given a five-minute period for practice browsing outside the SI Art Image Browser. Only one nonexpert subject required use of the practice time.

Prior to the first task, each subject was asked a few questions in order to collect demographic and control data. The research team randomly assigned each subject to a combination of tasks and search modes for each iteration of the test. Subjects were assigned to two of the SI Art Image Browser's three different search modes: (1) browsing for images using a classification hierarchy; (2) direct searching for images using keywords; and (3) a hybrid combination offering the subject equal access to both search modes (1) and (2). We balanced the search mode assignments for the first eighteen tests in each group to ensure a minimum of six subjects per search mode.   Subjects were not informed in advance of the system mode they were using.

All subjects were asked to complete two tasks during the testing period. Sample tasks included: "Find an image of geese flying," "Find an image of a bridge by Hiroshige," or "Find an earring and necklace set." The research team selected a total of fifteen test tasks in order to make available a variety of searching tasks. The tasks selected covered a range of search attributes including author, subject, medium, and combinations of these and other attributes. Team members pre-tested each task to ascertain its level of difficulty and mapped each task to an answer set of one or more images. For each task included in the test, the correct image or images could be located by any one of a number of possible search paths.

Upon completion of each search task, the subjects were asked a number of post-task questions to determine how familiar they were with the kinds of images they were searching for and to determine whether the subjects believed they had successfully found the correct image(s) to complete their assigned task. The test administrators also questioned participants about their experience, focusing on the usefulness of the system tested, any difficulties encountered, and suggestions for system improvements. After the subjects had completed both tasks, they were asked to compare their experiences using the two systems. These questions focused on general issues regarding the subject's use of the two systems and comparisons between the two. The entire test administration procedure took approximately one hour.

The data we collected included information relating to 20 variables: level of art domain expertise; gender; age; frequency of computer use; familiarity with Web; prior exposure to the Art Image Browser; practice with Netscape Browser; searching exercise assigned; browser mode (search, browse or hybrid); time to complete task; success in locating image; quality of search strategy; overall performance; prior knowledge of task's subject content; a self report of the subject's performance; percent of time system was believed to be useful; perception of utility of textual information attached to image; perceived differences in helpfulness between the two systems used; satisfaction with system(s); and preference for using the system again.

Profile of subject population
Two-thirds of the 31 subjects, 9 experts and 22 nonexperts, were female. Two-thirds of expert subjects were between the ages of 26 and 40, while nonexperts were younger - slightly under half in this range and about a third under 26.

The group had considerable familiarity with using computers and the World Wide Web. All expert subjects reported daily computer use while 91% of the nonexperts reported daily computer use, and the remainder reported weekly use. Over half of the expert subjects reported using the World Wide at least weekly, and 33% monthly. Over three-quarters of the nonexpert group reported using the Web at least weekly, 24% only a few times per year, with 9% never having used the Web before.

Many of the domain experts who participated in the controlled study had prior experience consulting various art-related databases useful to their own work and research. Over three-quarters of the expert group had had some prior use of the SI Art Image Browser, while none of the nonexperts reported prior use. As to be expected, experts brought to their task some familiarity with the subject content as well. Upon completion of their tasks, all nine experts reported being familiar to some extent with the kinds of images they were looking for in their

search tasks. However, the images were known to a sizable number of the nonexpert group: 16 out of 22 of the nonexpert subjects said they were familiar to some extent with their images.

Results

Does a browsing approach work?

One set of our research questions had asked: *Is a browsing model, in which images are classified into navigable hierarchies, a useful strategy for providing intellectual access to digitized images? How well would users be able to locate specific images or groups of images using this system? Are there performance differences between the three search modes of browse, search, and a combination of the two?*

In general, expert users were able to complete their tasks quickly and with a high degree of success. The nonexperts took somewhat longer to complete their tasks, experienced a lower degree of success overall, and had somewhat greater difficulty completing their tasks. All expert subjects reported that they had succeeded in finding images to complete their assigned tasks. Nonexpert subjects, while not as successful, were still able to report a success rate 86% of the time. We also found that no significant difference in terms of success rate appears likely between the three search modes the subjects used.

Of the 31 subjects questioned, 22 believed that the system was helpful 100% of the time, while only 2 thought the system was only helpful 25 % of the time, and no subjects found the system not useful at all.

The research team conducted a failure analysis of the test results which revealed that nonexpert users were more likely to fail to complete their tasks than were domain experts. Nonexperts also experienced greater difficulty with the direct search system than they did with the browsing system. The hybrid system helped both classes of users by providing them with alternate modes of access to the image collection if they should have difficulty completing a task by either browsing or searching alone. The greater portion of the failed tasks could be attributed to user confusion about how to use the search interface, which may be overly long and complex in its design. Compounding the problems users had with the search interface was user difficulty in operationalizing tasks into viable search terms. Finally, the research team found that task failure could also result from browsing categories which were too large or insufficiently subdivided, as excessive numbers of thumbnail images in any given category tended to discourage users from browsing.

The SI Art Image Browser was well received as an image retrieval tool by both domain experts and nonexperts. Comments by nonexpert users about their overall experience were very favorable. Typical of these comments were: "I enjoyed the experience," "very thorough system-- inviting to use," and "I like it a lot; very well put together, professional."

How feasible are thumbnail-size images for browsing?

To what extent can we rely on visual image information alone in browsing images? What types of accompanying textual information do users need in their search for images?

The subjects reported that the textual descriptions the systems provided were either very useful or somewhat useful in 30 of the 31 cases. Both groups favored higher resolution images, where this could be provided without seriously affecting the image loading time. Domain experts generally would have preferred higher resolution thumbnails so that the images could be better discerned, although one expert described the thumbnails as "fun, enjoyable, entertaining and beautiful" as they were. However, the experts also tended not to want to rely too much on the images alone; rather, they wanted descriptive information about the images to help them make their choices. Nonexperts were much more likely to select thumbnail images on a purely visual basis. They relied more on the images themselves and not so much on the textual descriptions. For example, one nonexpert subject described his selection method as a process of scanning over a page of thumbnails looking for "something different" or "something that stands out."

The nonexperts concern with the thumbnails' visual utility put them at odds with the expert users, whose comments indicated that they were more concerned by the loss of scale information engendered by the variations in thumbnail sizes in the test system. Some of the subjects noted that the size of the thumbnails made it hard to distinguish what was being depicted in the image. In these cases, subjects tended to want thumbnail sizes to vary depending on how clear the content of the image was.

Will a structured browsing approach appeal to a broad set of users?

Will a structured browsing approach be useful for a broad array of users interested in art images - both generalist users as well as subject domain specialists? Will a classification approach be useful to expert users, who would be more likely to have the domain expertise to search directly?

When presented with a choice, domain experts tended to use the search interface if they knew what they were after. They found the search categories to be logical. Many appreciated the internal consistency of having the search categories mirror the top-level choice of categories in the browsing system. A few technical problems with artist name searching became apparent during the domain expert testing, arising from problems in inverted rather than direct name order, misspellings, and accented characters in artists' names   Subjects were not always aware that they should include accented characters in their searches, nor were they necessarily aware of how to go about typing in diacritics using a standard keyboard.

The browsing mode dealt much more effectively with the artist name difficulties by presenting the names in the form in which they would appear in the database, that is, with diacritics. The domain experts generally believed that direct searching was quicker, more efficient, more "controlled," and allowed the user to "zero in more quickly" on what he or she wanted. However, the search mode was not seen to be without inherent difficulties, notably that a null search return could leave the user without an alternative plan for searching the database. Expert users' comments on the browsing system were very favorable, even though they tended to prefer searching over browsing. They felt that the browsing categories were logically organized and that the browser's classification structure and straightforward hierarchy made it easy to use.

Nonexperts seemed to find the search interface a little more frustrating to use than did the expert subjects. Nonexperts tended to form search strategies based on extracting as many keywords as possible from their assigned task and plugging them into the search interface. Although some

enjoyed the challenge of seeing whether or not they could get a result from the search interface, most were left perplexed if their best attempt at a search yielded zero hits. Some instinctively understood the logic of broadening a too-narrow search in the face of a null search return. Others simply shuffled the same set of keywords among different search fields. Their experience led to many constructive suggestions for how the search mode might be made more tractable to naive users. One suggestion was to provide a "search all fields" option that would eliminate the problem of user error in the selection of search fields. Other nonexperts suggested such features as options for "exact match," "close match," and "loose match," as well as Boolean and proximity operators. Some nonexpert subjects expressed an interest in having the system provide an online thesaurus of terms or pop-up menus of terms for each search field. Finally, these users suggested that a results ranking scheme would be very useful to them in making their selections.

In comparing the systems, the ratios of preference for experts were equal when comparing the search to hybrid system; 2:1 when comparing the hybrid to the browse systems; and 5:1 when comparing the search to the browse system. The ratios of preference for nonexperts were 5:6 when comparing the search to the hybrid system; 2:3 when comparing the hybrid to the browse system; and 7:5 when comparing the search to the browse system.

Will one set of categories fit all?
Can a single set of categories serve the needs of art history experts as well as generalist users?

Although the textual fields used in the SI Art Image Browser had been designed with the conventions of art historical description in mind, they appeared to satisfy the basic needs of nonexpert subjects as well. Nonexpert users repeatedly expressed a desire to access information about the artistic period or style of the images they were seeing. There was a clear sense, for instance, that these users would have liked to be able to search for art works broadly categorized as "Impressionist" or "Ming Dynasty" pieces. In the early development of the prototype system, the research team members had considered a field for precisely this kind of information but had abandoned the idea due to opposition from consulting art historians and curators, who presented a number of arguments against it.

The terms "object type" and "medium" were less intuitive to the nonexpert subjects than were the other categories. Subjects frequently cited "object type" and "medium" as field names which had caused them some confusion, whether in browse mode or search mode. The browse mode provided some guidance in explaining the distinction, with examples, but the search mode did not allow this comparison. Experts made additional comments about the text that bear noting here, including a suggestion to rename the generally confusing "object type" and "medium" fields to "medium" and "materials," respectively. This terminology would have to be tested with expert and general users to see whether it would indeed provoke less confusion about the content of these fields.

Since many of the experts did have prior experience in using various kinds of search interfaces in their work and study, most of them found the system expedient and straightforward; some suggested that they would not change anything about the design and function of the test system. Although the choice of descriptive data had been selected largely with the conventions of expert

users in mind, some expert users felt that the database was directed at a fairly basic level that would be interesting not only to experts but also to the greater number of general users.

Nonexpert users tended not to prefer long lists of choices in the browse mode; rather, they felt the long lists of terms indicated the need for further subdivision in the classification structure, and suggested increased classification subdivisions that would minimize displaying long lists of terms for users to scroll through. In addition, nonexperts would like to know in advance how many thumbnails they can expect to have load on a given page when using the browser, so they could not determine in advance whether or not it was worth the time to seek images there. Some felt that the search interface could be made more user-friendly if it were designed to include pull-down menus of terms used in the database. However, this would potentially create another problem of confusion where the pull-down menus of terms were excessively long.

Metadata for image searching
What types of information do users find useful in an image search?

Even though they tended to be more visually-oriented when examining pages of thumbnails, generalist users echoed the experts in their preference for thumbnails accompanied by very brief descriptions, such as artist, title, and date information. Experts advocated generous amounts of descriptive information about the images, including the dimensions of the original object. Most of the domain experts were in agreement that the system would benefit from additional contextual information, such as artist biographies, art criticism, bibliographies, etc. The experts felt this kind of contextualizing information would be especially important if an art image database were being used for educational purposes. Furthermore, they expressed an interest in being able to locate images by aspects of the image content, such as color and shape. Nonexperts echoed the experts' desire to see more than the basic art historical descriptors of the images, such as artist, title, date, etc. They would like to have seen more contextual information, such as artists' biographies, other works by same artist, links to groups of similar objects, even historical information or art history criticism. Nonexperts also expressed an interest in searching images by other means than those which were available to them in the test system, such as holding institution, acquisition date, or (like the expert group) by image content such as color, shape, or style.

The text fields selected for display with the images had been developed in consultation with individuals experienced in using and maintaining art museum databases. Consequently, most of the domain experts tested appeared to be at ease with the choice of fields and the kind of content these fields provided. When questioned about possible improvements to the images' textual descriptions, one expert responded, "I can't think of any, basically because I work and think using the same categories myself on a daily basis." Nevertheless, there did appear to be one field which the experts found noticeably absent and that was "dimensions." Many experts felt that the digital format of the images, along with the varying sizes of thumbnails, contributed to a loss of scale, which would have been countered, at least to some extent, by the provision of information about object dimensions. The team members had considered including dimensions at an early stage in the development of the test system. This information was ultimately left out due to the difficulties involved in extracting the information from existing museum databases, which record the dimensions in multiple fields and in differing ways depending on the physical nature of a given object.

Nonexperts tended to view the images as artifacts in and of themselves rather than strictly regarding them as surrogates. As such, the nonexperts were less concerned than domain experts with object dimensions and with the loss of scale engendered by differing sizes of thumbnail images. They were also less dependent on textual descriptions for identifying interesting or relevant images, especially where the textual data was unfamiliar or confusing to them.

Both experts and nonexperts suggested that incomplete or inaccurate descriptive data would compromise the functionality of either a searching or browsing system. Although the nonexpert subjects were less inclined to seize upon inconsistencies in the database, they also found missing data a source of confusion. Often, the problem of missing data was that it could mislead subjects and create difficulty in interpreting the results of their searches. This was especially the case with subject indexing. Only a portion of the total number of images in the prototype system included subject terms as part of their descriptive data. This was due in large part to the cataloging practices of the image providers who contributed images and data to the SI Art Image Browser. In many cases, subject terms were not used in-house at all by the image providers since subject terms may be considered, like stylistic period information, somewhat superfluous knowledge to art domain experts. Where image providers did include subject terms in their art object cataloging, the choice of terms could be inconsistent.

Expert users were quite sensitive to perceived inconsistencies in the data, particularly in terms of inconsistencies in the way in which art objects had been cataloged by their holding institutions. At least one domain expert pointed out that "dirty" data can greatly impede the function of an otherwise capable system. Many of the expert users tested pointed out, in particular, that subject terms should be present for all images in order to make subject searching an effective strategy, and that subject terms should be present in the correct form (for example, a junk is not the same as a sailboat). In large part, the users didn't realize that they should assume there would be missing data for many objects and that they should plan their strategies accordingly. Interestingly enough, this expectation that object data would be complete prevailed in spite of the experts' own experience with incomplete data in the art object databases with which they were most familiar.

Summary and conclusions

The SI Art Image Browser research project utilized broad classification categories to support browsing, on the premise that browsing can serve an important role as a strategy for retrieving image information. Since browsing relies on the user's cognitive abilities to recognize items of interest without having to formulate a specific query, it can be better suited for generalist users in an image search. Bawden (1993, p. 72) characterizes browsing as "[f]ollowing a predefined characterization, probably hierarchical, through an information space, to identify interesting items; [o]btaining an overview of the variation in items in an information space, so as to identify interesting areas of that space." Classification can play a key role in providing structures to assist users in navigating sets of images, and has the potential to assist users looking for specific images or information as well as help those with less well defined queries, and where key concepts were difficult to operationalize.

While this strategy will never be a substitute for the type of access made possible by exhaustive indexing, the results of the Art Image Browser study indicate that such a system can be effective in providing access to a collection when it is not possible or feasible to provide conventional levels of indexing. It also suggests that browsing can be an effective way for generalist users to gain intellectual access to images in digital collections without having an extensive knowledge of terms and concepts in the subject domain of the collection. Results of the controlled study showed no significant difference between the three search modes - browse, search, and combination of the two, used by the two groups of specialist and generalist subjects. If browsing proves effective as a strategy for retrieving images, then the development of classification to assist users in structuring their browsing should be a rich opportunity space.

The approach utilized in the SI Art Image Browser can have wide application to the management of image information in a diverse spectrum of disciplines and some members of the research team have used this model to build a retrieval system for images in earth and space sciences at the University of Michigan (http://focus.eecs.umich.edu/). The strategy can serve as an alternative access when detailed indexing is too costly and time-consuming and can also complement access when detailed indexing has already been provided for visual images.

Acknowledgements

References

Arnheim, R. (1970). Words in their place.The Journal of Typographic Research , 4, 199-212.

Art & Architecture Thesaurus., 2nd ed. (1994) New York : Oxford University Press.

Bates, M. J. 1989. The design of browsing and berrypicking techniques for the online search interface. *Online Review* 13, 407-24.

Bawden D (1993) Browsing: Theory and practice. *Perspectives in Information Management*, 3, 71-85.

Besser, H. (1990). Visual access to visual images: The UC Berkeley Image Database Project. *Library Trends* 38, 787-798.

Besser, H. (1992). Adding an image database to an existing library and computer environment: Design and technical considerations. In Susan Stone and Michael Buckland (Eds.), *Studies in multimedia: State-of-the-art solutions in multimedia and hypertext* (pp. 31-45). Medford, NJ: Learned Information.

Cawkell, A. E. (1992). Selected aspects of image processing and management: Review and future prospects. *Journal of Information Science* 18, 179-192.

Cawkell, A. E. (1993). Picture-queries and picture databases. *Journal of Information Science 19*, 409-423.

Chang, Shan-Ju and Ronald E. Rice (1993). Browsing: A multidimensional framework. *Annual Review of Information Science and Technology* 28, 231-278.

Enser P. G.B. (1995). Pictorial information retrieval: Issues related to subject indexing *Journal of Documentation 5*, 126-70.

Ester, M. (1990). Image quality and viewer perception. *Digital Image-Digital Cinema*, Supplemental Issue, Leonardo, 51-63.

Grosky, W. I, et al. (1994) Using metadata for the intelligent browsing of structured media objects. *Sigmod Record 23*, 49-56.

Grosky, W. I. (1997). Managing multimedia information in database systems. *Communications of the ACM 40*, 73-80.

Hastings S K (1997) A comparison of retrieval problems for digital images in a distributed network versus a closed system. In: Hemsley, J, (Ed.) EVA'97 Paris: Images Numeriques Appliqués aus Arts Visuels.

Hastings, S. K. (1994). An exploratory study of intellectual access to digitized art images. Ph.D. Dissertation, School of Library and Information Studies, Florida State University.

Holt, B. and Hartwick, L. (1994). Retrieving art Images by image content: The UC Davis QBIC Project. *Aslib Proceedings 46*, 243-248.

Kwasnik, B. H. (1992) The functional components of browsing. *Annual Review of OCLC Research,*, 53-56.

Shatford Layne, S. (1994). Some issues in the indexing of images. *Journal of the American Society for Information Science 45*, 583-588.

Markey, K. (1988). Access to iconographical research collections. Library Trends, 37:154-174.

Stam D C (1989)Tracking art historians: on information needs and information-seeking behavior. Art Libraries Journal, 14:13-16.

Appendix

Questions used for image search

Task
1   - Find the works by Kandinsky
2   - Find the works by Manet
3   - Find an image of Dostoevsky by Beckman
4   - Find a sculpture by Picasso
5   - Find an oil painting by Renoir
6   - Find an image of a bridge by Hiroshige
7   - Find a mask made of shell
8   - Find a mural that contains Christian religious figures
9   - Find a Japanese woodblock print that contains a sailboat
10 - Find a stoneware vase
11 - Find an image from the Koran
12 - Find an image of a mother and child
13 - Find an image of geese flying
14 - Find an image of a reclining nude
15 - Find an earring and necklace set