Date submitted: 02/06/2010

Words … will not stay in place: cataloging and sharing image collections

James Shulman
President, ARTstor[1]
151 East 61st Street
New York, NY 10065
USA

Meeting: **71. Art Libraries**

WORLD LIBRARY AND INFORMATION CONGRESS: 76TH IFLA GENERAL CONFERENCE AND ASSEMBLY
10-15 August 2010, Gothenburg, Sweden
http://www.ifla.org/en/ifla76

Abstract:

Words have been affixed to still images for hundreds of years to describe who created a work or what it portrays. This paper examines the ways that this process might evolve in an era where dissemination of knowledge is far less linear than it was in an age of print. The paper reviews two projects that ARTstor is pursuing with the art historical community. One captures users' notes about 190,000 photographs of old master drawings that are in need of updated descriptive cataloging. The other effort will create a Built Works Registry through the use of a credentialed Wikipedia-like strategy.

…Only by the form, the pattern,
Can words or music reach
The stillness, as a Chinese jar still
Moves perpetually in its stillness…

…Words strain,
Crack and sometime creak, under the burden,
Under the tension, slip, slide, perish,
Decay with imprecision, will not stay in place,
Will not stay still.

T.S. Eliot *Burnt Norton* V:140-142, 149ff

[1] This paper draws heavily upon the ideas, prose, and suggestions of current and former colleagues at ARTstor including Carole Ann Fabian, Christine Kuan, Max Marmor, Kathleen Schowalter, Gretchen Wagner, Dustin Wees, and Bill Ying.

In *Burnt Norton,* T. S. Eliot struggled with the question of whether words could be depended upon to convey meaning, whether they could capture a moment of vision and then consistently convey that vision to others. While, as a poet, he sought to choose words that could capture a feeling, he also recognized that every attempt to have a word reify a moment depends on a constancy of meaning for which there is no guarantee. Whatever meaning the word had sought to capture would move "perpetually in its stillness."

Words have been affixed to still images for hundreds of years to describe who created the work or what is portrayed in the work. Such descriptive cataloging not only allows visual media (such as images) to be discovered via text based searching (either in file cabinets or in databases) but also conveys information that is not intrinsic to the image (such as who created the work). Yet that assignment of words – like Eliot's poetic words – is also contingent, the product of a moment of choice for a scholar or a cataloger acting at a certain point in time, in a particular institutional and personal context, and working within a certain consensus about the work. But knowledge of the work might change: many works once labeled as created by "Rembrandt" are now no longer considered the work of the artist. In this paper, we will examine the ways that the process for affixing words to images (and thus making them discoverable and meaningful) might evolve in an era where dissemination of knowledge is far less linear than it was in an age of print.

We will look at the structure, aims, and methods of two projects that ARTstor is working on with the art historical community.[2] One of these projects involves users as builders (rather than passive recipients) of an extraordinary collection of images of old master drawings that is in need of updated data. The other effort will create a powerful cataloguing utility for architecture, landscape, and the built environment through the use a Wikipedia-like strategy. Finally, we will close with a brief exploration of how we will support these efforts in the evolving ecosystem of digital scholarship.

[2] ARTstor, a nonprofit organization, was founded by The Andrew W. Mellon Foundation, with a mission to use digital technology to enhance scholarship, teaching, and learning in the arts and associated fields. As of June 2010, ARTstor provides software, services, and 1.1 million images from over 190 source collections to over 1,250 subscribing institutions, and is developing (in partnership with colleges, universities, and scholarly societies) a networked cataloging and image management system (called "Shared Shelf.")

Background: the linear model for attaching words to images

For centuries, images of art works have been cataloged—described either by an expert, librarian, or curator at the holding institution or by an outside scholar studying the image. Sometimes collections of images had their own life as the object of cataloging, with the cataloging being done by teams of scholars. Consider, for example, the multi-volume catalogue of old master prints originally created by Adam von Bartsch and further developed in the 20th century as a 96-volume illustrated catalogue under the direction of Walter Strauss and involving dozens of scholars, or, the *Corpus Vasorum Antiquorum*, created originally at the Union Académique Internationale of France in 1919 and eventually branching out to include 100,000 vases located in 26 countries. We are the beneficiaries of the manifold photo archives created by organization or an individual with a discerning eye left in the charge of a scholarly library to be preserved and made available. The numbers of these collections are countless–from the Berenson Photo Archive, the Foto Reali Archive (National Gallery of Art, Washington, DC), the Robert van Nice archive at Dunbarton Oaks, the A. C. Cooper Archives (Frick Art Reference Library), and countless others.

For these collections, the techniques of librarianship were applied in practical ways; this meant, in many cases, creating finding aids to organize material that could not immediately be cataloged at the item level. It meant bringing not only the tools of scholarship to focus on the cataloging of particular works but also the tools of library science so that works could be organized in taxonomies and terms could be drawn from controlled lists or authority files so that the works would not disappear in the sea of noise that results from idiosyncratic cataloging. These great efforts toward organization, preservation, and access have been carried out with painstaking effort, significant institutional and personal investment, and long-term vision. As with other legendary, intergenerational projects conducted by the library world, such the Middle English Dictionary which took 75 years to complete, this sort of heroic undertaking enabled what might be thought of as a linear progression of content provision. Work was done at one place (perhaps at the site of the works) or by a circumscribed team of experts; their work, duly edited, was then distributed to scholars who would visit the edited resource or subscribe to copies of it. Mostly, the products of

cataloging were images that found their way into various publications: collections catalogues, exhibition catalogues, scholarly journals or books. The work still lived in a place, the image of it may have been more widely circulated, and the interpretation of that work–the words attached to the image might live anywhere from the repository that owned the work to a printed article or book. Somewhere between archival management and publication, the products derived from primary source visual material emanated outwards from their source.

The Gernsheim Photographic Corpus of Drawings

The Gernsheim Corpus is an extraordinary scholarly resource: a professionally-produced photographic archive, consisting of approximately 190,000 photographs of European old master drawings from over a hundred European and American public and private collections. Built up over the course of a seventy years by Swiss photographer Walter Gernsheim and his wife and collaborator, Jutta Lauke Gernsheim, the Gernsheim Corpus can only be consulted in its entirety in a half dozen subscribing institutions in Europe and the U.S. When we first met in 2003, the then 82 year old Mrs. Gernsheim and her daughter Karin had just spent two weeks in Palermo shooting; she then shot in Munich where she had been waiting to shoot for 50 years, where internal dynamics within the institution had now changed in her favor. She spoke eagerly of plans to shoot in Poland, Budapest, and the Hermitage. When they were on a campaign, they would work twelve-hour days, taking many shots of each work, and then return to the hotel room near the museum to develop the film in the bathtub so that they were able to review the photographs while their memory of the works was still fresh. It is a simply astounding resource.

The Gernsheims saw the nature and purpose of ARTstor as consonant with the original purpose of the Corpus: to use the best technology available to promote access to materials necessary for non-commercial scholarly research. While they felt comfortable providing us with the entire corpus to digitize, they also felt that ARTstor should be working closely with the holding museums in incorporating the Corpus into the ARTstor Digital Library.

But the current state of knowledge about the works portrayed in the photographs records has not stayed still; the words, compiled over the course of 75 years, need updating. At a meeting of the drawings keepers in 2005, we learned what we already knew (on some level): attributions of drawings that were made as early as 75 years ago had often been either supplanted by a different attribution or were (to the mind of the curator) in need of new research. Since the model that curators knew for dissemination of these images was publication–reproduction in a fixed and printed journal or book, the reluctance of a few of these curators to be enthusiastic about the release of these images for use in ARTstor carrying their old attributions was completely understandable.

But the old model of print publication in journals or books was not the only way that these images were consulted or discussed. When we first started working with photoarchives at the Frick and the Getty in 2001, we learned that they were adding to the cataloging of Gernsheim records (on the basis of their own staff efforts and the efforts of visiting scholars), but these additions were locked within the physical barriers of manila folders and institutional walls. Although information scribbled on the mount of a photograph in the photo archive was available to the next scholar who looked at that image in that archive, and thus more dialogic than the still publication frozen in its binding, these notes were not circulated between archives or returned to the institution that cared for the original work.. The spaces were scholarly workspaces where engagement with an image also meant engagement in a dialogue–not only with one's colleagues of today but with those who had come before. This sort of intergenerational dialogue took place at a pace far different than "real time." Still, despite some familial resemblance to the keyword tagging that users add today to images in Flickr or YouTube, the addition of attributions conducted by staff or scholars in libraries leans a little further toward the conservative side of the authority scale than the wide-open contribution model often employed on the open Web. The Web's transformation into Web 2.0 came–for ARTstor–at the same time that we began to build a critical mass of both images and users. Without either, it would be presumptuous to ponder whether there was a better workspace for enriching and enhancing the dated and centralized Gernsheim corpus; but, today, we believe that there is.

ARTstor has a plan to create the digital and networked workspace that the print publication model could not support, by carrying out a responsible and, above all, an achievable solution that addresses the most significant curatorial concerns about the project without imposing insupportable burdens on ARTstor and the museum community. To this end we will: 1) encourage ARTstor users to recognize that the data provided with the Gernsheim Corpus are to be regarded as provisional rather than definitive; 2) assist ARTstor users by directing them to more current scholarship on the drawings collections represented in the Corpus, and 3) create a workspace that allows individual, identified users to associate cataloging, commentary, and bibliographical links with an individual image–data which will then be available to the community and also to the individual curator at the institution charged with authorizing the data about any given work.

Briefly, here are the steps to enfold the Gernsheim Corpus into a community-based workspace – where authority still matters but is developed over time and in dialogue between others who care about the works.

Step 1: Make the Gernsheim data and images available

The baseline dataset for the Gernsheim Corpus comes from the original inventory lists provided by the Corpus to its handful of subscribers. The Gernsheim inventory is not decipherable to readers as it includes "codes" for media that cannot be programmatically changed to English words due to the fact that the codes have different meanings within the Gernsheim data (e.g. "W" indicates both "white chalk" and "watercolor"). To that end ARTstor will attribute ARTstor's "original data" to the Gernsheim Photographic Corpus of Drawings[3] by: linking creator names to the Getty Thesaurus Union List of Artist Names (ULAN); cleaning the data in the creator field data so that the scholars' names who attributed these works and other shorthand notes are moved to the notes field in the data records; and converting Gernsheim material and medium codes to English words.

[3] Each record includes this statement: "Descriptive data, especially artist attribution(s), may have changed since the descriptive data was collected by the Gernsheim Corpus. For more current scholarly opinion, ARTstor users are encouraged to consult the scholarly literature and especially the owning repository's publications and website."

Step 2: Direct users to (or provide them with) more current data

In every record, we will include a URL to the repository that holds that drawing; we also will reach out to participating repositories and offer to incorporate updated information supplied by the repository. We have already done this with the images launched from the British Museum collection.

Step 3: Enable individual users to contribute to the conversation around the work

Each image file in ARTstor has an accompanying data record that contains whatever descriptive cataloging we received from the source, data about the size and source of the image file, and space for individual users to make notes (either for their own personal use or for a group/class of users on their particular campus). We intend to create a new tab in our metadata window called "Bibliography/Discussion" that would be open only to those who had registered as "instructors." By entering one's name and affiliation, one would be allowed to provide links and comments that would then be associated with that image for all ARTstor users at all participating institutions to see. As links and comments compiled, the "conversation" around a work could become a significant locus for community building among teachers and scholars.[4] As the photoarchives of the past 100 years have shown, the individual image is the appropriate locus for this community-built referencing. From the ARTstor digital library's perspective, we could never hope to suggest bibliographical links on our own (since any attempt to do so algorithmically would be at an exceedingly high level of generality–i.e. at the level of the artist rather than the specific work at hand) and hence not very useful. But the discussion at the level of the work, furthered by those who are focused on that particular work, could–if enough scholars are engaged

[4] We would also seek to make these links and comments compliant with a new emerging standard being developed by the "Open annotation" project. This effort seeks "to facilitate the emergence of a Web and Resource-centric interoperable annotation environment that allows leveraging annotations across the boundaries of annotation clients, annotation servers, and content collections." By having user-added links and comments to be managed in a standards-driven way, we could have greater comfort that the investments made by individual contributors will live on as the content migrates (both within ARTstor and, potentially, outside of ARTstor eventually as well)

via a broad enough network–provide a highly scalable version of what once was a local activity constrained by the physical presence of the printed photograph.[5]

A curator may or may not choose to delve into the discussions that occur around a work that their institution holds. In a growing array of communities, participation from the passionate amateur is welcomed by the professional community (and, of course, since ARTstor only reaches educational and cultural institutions, the amateurs who would be involved would be drawn from college and university instructors and researchers as well as fellow curators). Consider, for example the Center for Backyard Astrophysics, "A global network of small telescopes dedicated to photometry of cataclysmic variables:"

> More than half of our time is currently devoted to the study of "superhumps" -- light variations at periods near but not exactly at the binary orbital period (typically a few hours). Superhumps are ... very difficult for professional astronomers to study, because they're often transient, and because they're just not very well constrained by observations over a short time baseline. Backyard telescopes solve both problems. We find that by amassing large amounts of data over an observing season, especially with robotic telescopes and a range of terrestrial longitudes, we can build an observational record of far better quality. And we do.[6]

If the old Master Drawing community is able to draw upon observations and proposals of scholars that they know (at the 1,250 ARTstor subscribers and the 100+ Gernsheim contributing repositories), they are opening up a fairly wide dialogue. While the authorized attribution will continue to be the decision of the institution that cares for the work, there would be so much that we could do to recycle the discussion that accumulates around the work into the indexing. We could go on to identify (visually on the results page) images that had "Bibliography/discussion" associated with them and/or allow users to set a preference to rank results (or filter results) on the basis of whether images had such expert tagging. If the

[5] Currently only six institutions have complete sets of the Gernsheim photographs, drastically limiting the possible discussion of the drawings within the scholarly community. With our release of the collection and bibliography/discussion tools within the ARTstor digital library, our 1,250 subscribing institutions would have immediate access to both the full corpus of images and the scholarly discussion (an increase of more than 200 times/20700%). At the same time, the contributing Gernsheim repositories would get access to the digital Gernsheim collection and the discussion that ensued around the works).

[6] http://cbastro.org/background/history/

"Bibliography/discussion" section becomes a robust format for discussion, we could allow an advanced search function that would allow the user to search for comments by specific people or links to specific periodical literature. For example, a user might care a great deal about comments made by Keith Christenson or Elizabeth Cropper or Stefan Moret or articles that are written for *The Journal of the Warburg and Courtauld Institutes.* All of these enhancements could occur down the line, if the workspace earns a place for animated and active discussion; but if that does happen, we will have a great deal of room for borrowing innovations honed by the communities that surround Amazon, Netflix, E-Bay, and Wikipedia.

Our interest in including the Gernsheim Corpus in ARTstor was first piqued in 2001 when we saw a project that the Frick Art Research Library was doing building a database of Italian anonymous works. Those works labeled with the contingent word, "anonymous" were, of course, created by people just as the attributed works were. The networked approach discussed here could someday change those nameless works to named parts of our collective cultural heritage. The de-centered networked environment can allows us to have those conversations, and provide some contingent answers.

Developing a registry for the built environment

If the effort to place the once static Gernsheim Photographic Corpus into a workspace wherein scholars can contribute to the words that will enrich access to (and knowledge of) an historic collection, a separate effort–a Built Works Registry–will allow knowledgeable individuals to build a resource to be employed in the creation of new records associated with works or architecture or landscapes. The goal of this project would be that the shared community knowledge of oft-photographed works, such as Frank Lloyd Wright's *Fallingwater,* could result in the beginnings of an authority file for the built environment. By calling upon such a file, any catalogue entry of a new image could build on previous cataloging, improving efficiency and—importantly—consistency.

Built works provide a particular challenge for catalogers of visual materials. Because there is no community-wide list of names of the built works themselves, but rather ongoing and improvisational

practices at individual institutions, based on idiosyncratic standards, or no control at all, cataloging can be both redundant and inconsistent, rendering processing and retrieval of records inefficient and sometimes even unusable.[7] For example, Wright's *Fallingwater* near Mill Run, Pennsylvania, is also known as the *Edgar J. Kaufmann Sr. Residence* or simply the *Kaufmann House*–often confused with the *Kaufmann Desert House*, a Richard Neutra-designed residence in Palm Springs, California. Adding to the difficulty of this challenge, built works can be complex, multi-part works without clearly identifiable creators or fixed dates. The BWR will help disambiguate one work from another by providing both a preferred set of data and allowing for variants.

The building of the Built Works Registry will be collaborative, engaging the national and international educational and cultural communities. The community has benefited tremendously from the authority file editorial efforts of institutions like the Library of Congress and the Getty Research Institute. But we believe that rather than placing the entire editorial burden on a single stress point (as projects like ULAN and LCNAF do), we should move towards developing collective models of contribution and editing. As the *Report of the Library of Congress Working Group on the Future of Bibliographic Control* notes:

> New partnerships will result from collaboration and coordination among a wide array of stakeholders. This will realize workflow efficiencies and minimize redundancies between and among entities that create and use both authority and bibliographic data... [8]

Although the Getty Research Institute's Vocabulary Program has outlined a schema for a work authority file, the *Cultural Object Name Authority* (CONA), there is no collective community plan to define policies, migrations, and shared infrastructure that will seed, build upon, and edit such a file.

---

[7] As the final report of an Andrew Mellon Foundation-sponsored project that sought to link records from numerous visual resources collections (the Union Catalogue of Art Images or UCAI) noted,

> Unique object identifiers would significantly improve processing and retrieval… For artworks and architectural structures, there is no equivalent to an ISBN or ISSN. …Architectural structures and sites have no such identifying system. An international and coordinated object identifier registry, perhaps modeled after the ISBN/ISSN, could provide an efficient method of identifying objects (and, presumably, works).
> http://gort.ucsd.edu/ucai/doc/ucai2_final_report.pdf, pages 7-8.

[8] http://www.loc.gov/bibliographic-future/news/lcwg-ontherecord-jan08-final.pdf, page 21.

Beginning work on such a file will be a very significant step forward in an undertaking that ultimately aims to map the entire built environment.

The Avery Architecture and Fine Arts Library at Columbia University and ARTstor would be the principal initial collaborators for this project. In addition, we are working closely with nine other partner institutions and organizations (Colby College, Cornell University, Harvard University, Middlebury College, New York University, the Society of Architectural Historians, the University of Illinois at Urbana-Champaign, University of Miami, and Yale University) on a networked cataloging and image management platform (known as "Shared Shelf") and we have also enlisted an experienced international advisory group.[9] This group will bring both wisdom and knowledge of others who have gone before us in this territory and have already accomplished a great deal.

To build the enabling infrastructure for such an undertaking is an incredibly complicated effort. For example, we will need rules about complex objects (are each of the houses at Seaside separate "works" or is Seaside the "work" or both?) On the interface side, how can one check a list for the name of the work since most built works are not named? Is the list sorted for the user by geography or style or period or architect (when there is one)? Needless to say, the software challenges (both the database issues and the interface requirements) are endlessly complex – and fascinating.

While the first stage of building such a registry is developing a schema and a set of policies, we also believe that the project will only begin to become real enough to matter if we can seed the file with records so that early users can see its potential utility. Seeding the file with lists of works from the ARTstor Digital Library, Harvard, and Cornell should result in a substantial starting file that is drawn

[9] Murtha Baca, Head, Digital Art History Access, Getty Research Institute; Barry Bergdoll, The Philip Johnson Chief Curator of Architecture and Design, The Museum of Modern Art; Tom Bilson, Head of Digital Media, The Courtauld Institute of Art; Jocelyn Gibbs, Associate director, Collection, Canadian Centre for Architecture; Monika Hagedorn-Saupe, Deputy Director Institute for Museum Research, Berlin; Dianne Harris, Associate Professor of Landscape Architecture, Architecture, Art History, and History, University of Illinois, Urbana-Champaign, President, Society of Architectural Historians; Irena Murray, Sir Banister Fletcher Director, British Architectural Library, Royal Institute of British Architects; Sam Quigley, Vice President for Collections Management, Imaging & Information Technology, Art Institute of Chicago; James Quo-Ping Lin, Research Director, Chief of Exhibition Service Division, National Palace Museum, Taipei, Taiwan; Gunther Waibel, Program Officer, Office of OCLC Research; Ann Whiteside, Head, Rotch Library, Massachusetts Institute of Technology; Hugh Wilburn, Director of the Frances Loeb Library; Assistant Dean for Information Services, Harvard University Graduate School of Design.

from image libraries built to serve educational users' needs. From such a list, policies about who can add new terms or add alternatives to existing entries will quickly be upon us, and we anticipate an enormous amount of work (and an enormous opportunity) if we are able to aggregate other lists from other countries and build tools to enable editors (be they ARTstor staff or community members) to merge existing terms into the Registry.

How we will support these projects

Harnessing the power of contributors and editors of a Built Works Registry can only happen if there is a software environment that allows those who want to help to do so. Enabling the sharing of enormous image collections and creating an online environment to allow for the exchange of scholarly ideas and issues around those images also requires a range of organizational support and technical maintenance. In his book, *We-Think,* Charles Leadbeater (an admirer of the wiki culture) notes that enfolding people's passion is not the only element required to unlocking the enormous power of collaborative work:

> The wiki-economy has not escaped the deep-seated problem that beset earlier attempts at collaborative endeavor. Communes, mutual societies and worker co-operatives often failed because they closed in on themselves and avoided hard decisions about how work should be organized and money made. [10]

ARTstor's work to digitize the Gernsheim corpus (and dozens of other important collections) was funded by The Andrew W. Mellon Foundation in the hopes that this significant expenditure could unlock very important "returns" for the scholarly and cultural community. ARTstor's content and its software environment are sustained by institutional participation fees which are scaled to accommodate a range of institutions.

Still, some elements of the system, such as the Built Works Registry, must be open to community contributions and community calls–without any fees at all. It would be foolish if we were to erect tolls for the willing cataloger, editor, or scholar to add to or improve such a registry. ARTstor and the Shared

[10] Charles Leadbeater, *We-Think; Mass Innovation, Not Mass Production.* London: Profile Books, 2008: p. 90.

Shelf software users will contribute to maintenance of the Built Works Registry technical infrastructure, but we will need to be mindful of Leadbeater's caution noted above. Warm feelings will not allow the infrastructure to keep pace with user demands.

The sustainability model that seeks to have all who benefit from a project pay in a fashion appropriate to their capacity and to the value that their organization will receive is not always popular when we all would prefer that such services would be free to all. But having such a mission-driven and community-minded support model will enable the support of some public goods (like the Built Works Registry). If the enabling infrastructure is valued and supported by the community that cares about the work getting done, the whole effort can be less dependent on the good will of contributors, foundations, or governments, and the service is more likely to continue to exist.

* * *

Widening levels of access in the building of new knowledge are the common element of these two projects. There will be some who may want to bar the gates against the noisy and messy masses. And if the gates were cast wide open, there might well be dangers in inviting in the "crowd." The crowd, in this case, surges with people who care passionately about the difference between Bandinelli's hand and that of Francesco Salviati, or who are willing to research the dates that Viipuri Library was built. We hypothesize that their contributions may contribute significantly to our shared cultural knowledge. Eliot's fear – that words will decay with imprecision – need not only be a source of anxiety. Thoughtful people will need to be able to choose new and different words to describe images and the network will allow a range of community experts to contribute their knowledge for consideration. Authority still has to be respected, consensus still has to be earned and assembled, and the circulation of commentary need not slide into the cacophony of internet chat rooms. At one time, authority flowed only from the pen of a solitary scholar to a printed volume, with chosen words frozen at a moment of time (via publication); in that approach, the contingent nature of the words chosen to describe an image was masked by the typeset

appearance of permanent truth.[11] We hope that these two experiments will demonstrate the potential of distributed knowledge rather than the undermining of all authority. The new consensus that a community of scholars assembles will, of course, also be imperfect and provisional; but, at least, the next "edition" will be only a few keystrokes away.

[11] In a similar way, many people often mistake the numbers in a spreadsheet for "facts," and fail to question the source or the validity of the data reported or even whether the formula masked beneath a "total" has added a column correctly.