HyperImage: Of Layers, Labels and Links

Heinz-Günter Kuper <kuper@inkubator.leuphana.de>, Jens-Martin Loebel <loebel@inkubator.leuphana.de>
Leuphana University Lüneburg

Abstract

HyperImage is a software platform geared towards assisting researchers in image-based disciplines. Since its initial development, the software has gained in sophistication and relevance to the digital humanities in general and to art history in particular.

The software allows users firstly to manage images and metadata from various external sources (e.g. image repositories) as well as from local storage media. Additionally, it supports the marking of regions in images and the linking of these regions to a diverse selection of link targets, be they other image regions, whole images, groups of images, text, etc. HyperImage thus visualises a web of discourse, simultaneously assisting scholars as a research tool and as a presentation tool. Semantic web technologies can be applied to the metadata and to the links between elements to permit researchers to mine the data contained in the system, while versions of a project can be made available to an audience for presentation, exploration and discovery.

This paper introduces a new version of HyperImage and explores its functionality within the context of the digital humanities, especially with regard to art history. A sober approach to the technologies involved will be presented, exploring not only their affordances, but also their limitations within the broader context of scholarly discourse.

Keywords

open source virtual research environment, digital humanities toolkit, annotation, linking, visualization

Introduction

Information technology has far more to offer the Humanities than databases, desktop publishing and presentation software. Technologies surrounding what has come to be known as the Semantic Web offer much potential to expand the horizons and toolsets of Humanities research. While the vision of the *macro* Semantic Web may never be fully realised, for example due to issues of semantic noise introduced by ever broader datasets or due to questions regarding the reliability of data (data provenance), *micro* semantic webs – i.e. datasets constructed around a particular focussed research question – offer modern Humanities researchers the opportunity to utilise information technology in the course of their investigations.

The focus to date has largely been on textual analysis, ignoring the image itself as the object of interest. The HyperImage Virtual Research Environment is a toolkit that redresses this imbalance by supporting the hyperlinking of (regions of) images and managing the complex metadata associated with this task.

In a HyperImage project, areas of an image can be defined as one or more polygons in a layer, thus allowing complex shapes to be defined. Each layer serves as a link anchor, allowing regions of an image to be hyperlinked to other content, such as other layers on the same or other images, text, and external Web sites. Additionally, a textual annotation may be defined for each layer. The textual annotation itself may also contain links. HyperImage is thus a tool that supports the second wave of the digital humanities, as elucidated by Schnapp and Presner in the *Digital Humanities Manifesto 2.0*:

> The second wave is *qualitative*, *interpretive*, *experiential*, *emotive*, *generative in character*. It harnesses digital toolkits in the service of the Humanities' core methodological strengths: attention to complexity, medium specificity, historical context, analytical depth, critique and interpretation. (Schnapp et al. 2009)

HyperImage supports image-based research not only by allowing the researcher to manage complex and detailed metadata, but also by enabling links to be created between images and other objects (such as external URLs), turning a HyperImage project into an information space that can be navigated in order to make new discoveries or to present the discoveries that have already been made.

Kuhn speaks of the changing approach to research that comes with a paradigm shift:

> During the transition period there will be a large but never complete overlap between the problems that can be solved by the old and by the new paradigm. But there will also be a decisive difference in the modes of solution. When the transition is complete, the profession will have changed its view of the field, its methods, and its goals. (Kuhn 1970, 84)

Kuhn goes on to quote the historian Herbert Butterfield, who describes the process as "handling the same bundle of data as before, but placing them in a new system of relations with one another by giving them a different framework" (Butterfield 1957, 1). The toolkits currently being constructed to serve the digital humanities are emblematic of this new framework, and HyperImage aims to provide art historians with novel access to old and new questions, leading to insights that would not have been possible within the paradigm of the traditional Humanities.

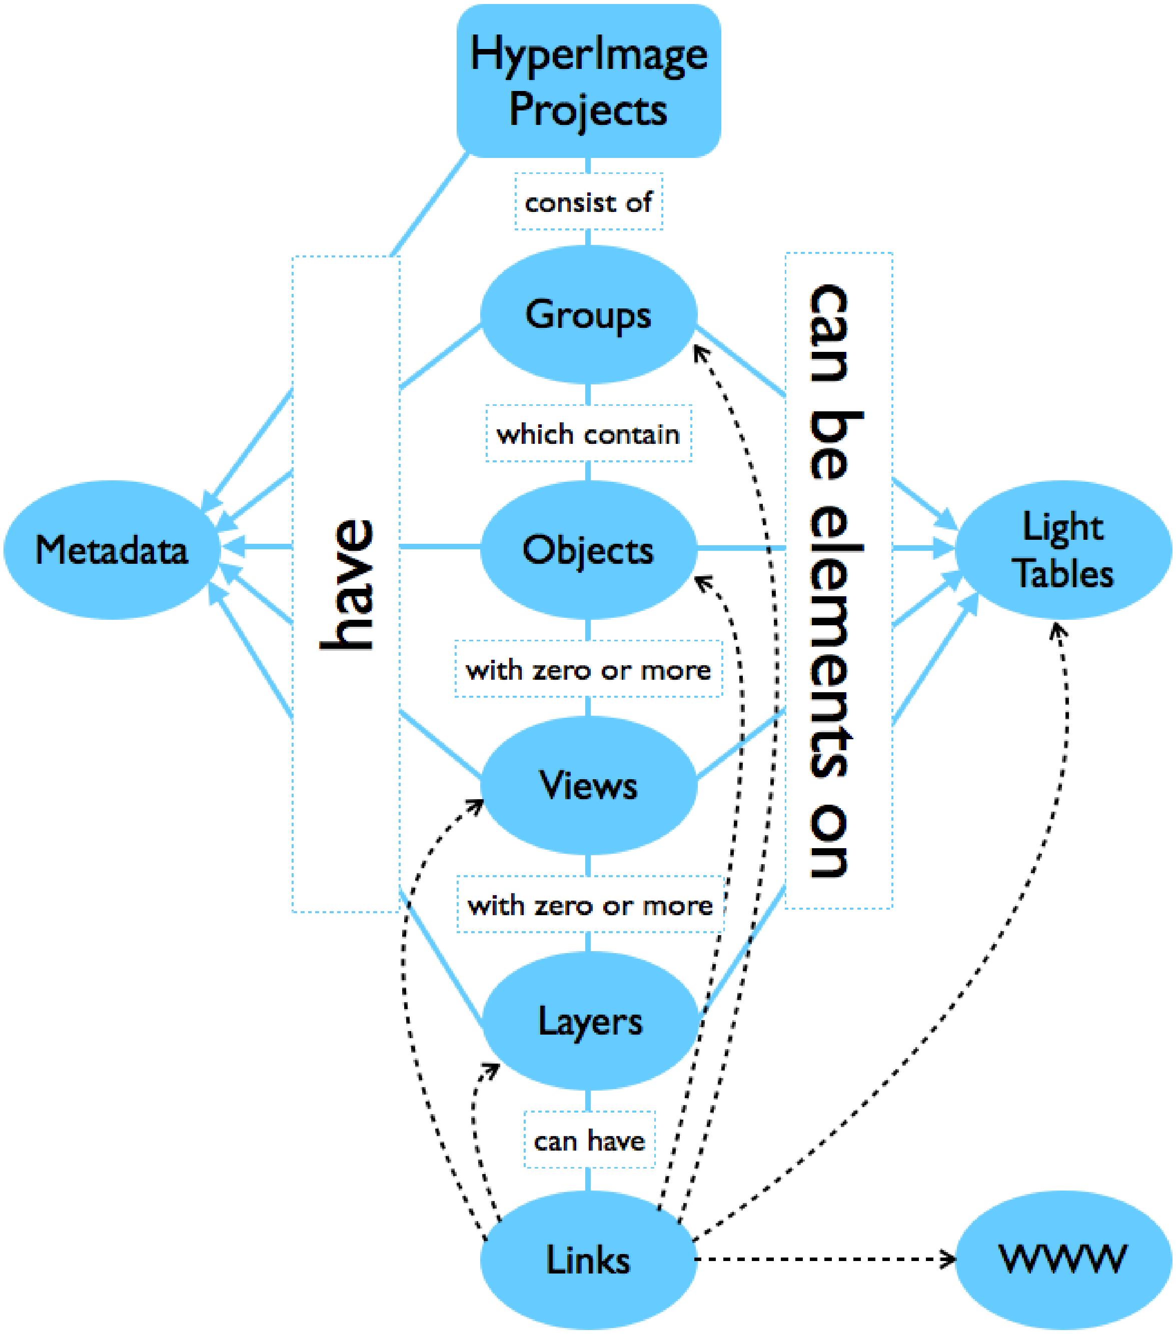

Figure 1: Structure of a HyperImage 2.0 project

A number of HyperImage projects have already been conducted in a wide range of fields, ranging from entomological projects at the Berlin Natural History Museum, to medieval prayer books, to street art, to propaganda images from the French Revolution, to zoological research, to virtual excursions through nature reserves, to visualising exquisite series of glass windows in medieval cathedrals. HyperImage is a versatile tool that can find application in many aspects of image-based research.

Typed Links and Semantic Web Technologies

A straightforward hyperlink between an anchor (that which is clicked) and a target (that which is displayed after the link has been clicked) can be considered to express the simple relationship "is_related_to", i.e. "anchor *is_related_to* target". Semantic Web technology, specifically the Resource Description Framework (W3C 2004), supports the expression of a multitude of relationships between anchors and targets, e.g. "is_author_of", "birthplace_of", etc. Using RDF it is possible to express typed hyperlinks, which in turn supports more sophisticated searching within a HyperImage project.

The Semantic Web can be summarised as enriching the World Wide Web as we know it with metadata. Internet search is currently based on keyword rankings: one inputs search terms into a browser and the search engine returns a list of pages where these terms are to be found, ranked according to an algorithm. A far more powerful form of search can be achieved by enriching the content *on* Web pages and the links *between* Web pages with metadata. For example, metadata might clarify that an artwork mentioned on a page refers to a piece of sculpture or disambiguate a place name when referring to Paris, Texas. Tim Berners-Lee's original conceptualisation of the WWW made use of metadata to typify links: e.g. a particular document "describes" some issue, "refers to" other documents, and was "written by" an author (Berners-Lee 1989). This would, however, have complicated the initial implementation of the WWW, so his vision was delayed. He revisited his idea in the seminal article "The Semantic Web" (Berners-Lee et al., 2001). Since then, the World Wide Web Consortium (W3C) has been pursuing Berners-Lee's vision and a number of tools and standards have been developed in pursuit of the goals of the Semantic Web.

Unfortunately, the Semantic Web in its grand conceptualisation of the "Giant Global Graph" (Berners-Lee 2007) is beset by two major problems, namely the issues of provenance and intractability. These issues can be traced to one of the maxims underlying the Semantic Web: "anyone can say anything about anything" (W3C 2002).

The data points spanning the Semantic Web can be thought of as nodes in a graph connected by logical predicates as edges. For example, the two nodes "Tim Berners-Lee" and "World Wide Web" could be connected by the edge "inventor_of", as in (Tim Berners-Lee, inventor_of, World Wide Web). Queries are resolved by so-called query engines and any query against the Semantic Web involves the traversing of a number of subgraphs that, taken together, form the Giant Global Graph. The computational processing involved is time-consuming and complex and unfortunately an optimal solution cannot be guaranteed.

The Web as we know it today is an open platform and, barring restrictive regimes and technical constraints, almost anyone can add content to the billions of Web pages already in existence. There is no guarantee that the information to be found on the Web today is trustworthy, but questions of provenance take on a new dimension when the steps in a deduction process executed by a query engine rely on dubious information. In other words, the user might construct a query and be presented with a result which is derived from several intermittent steps. How reliable is the final result when some of the intermediary steps are of doubtful provenance? Establishing the provenance of data contained in the Semantic Web is therefore very important, but not always possible.

In addition to the problem of reliability, there are logical constraints. Query engines process queries by applying logical operations and returning the results. Due to the openness of the Semantic Web, there is no guarantee that all the RDF statements involved in a particular query are non-contradictory. Such a query could never be resolved, either requiring the contradictory statements to be ignored or for the query to be made against a different dataset.

The Semantic Web at this stage is perhaps more of an idea than a concrete entity, but the conglomeration of datasets referred to as "Linked Open Data" is the closest incarnation of the idea we have to date. To see a list of the current datasets visit http://www.w3.org/wiki/TaskForces/CommunityProjects/LinkingOpenData/DataSets. While the LOD datasets have a strong Science bias, the focus is broadening as more and more datasets come online. Most notable among these is DBpedia, which exposes a subset of the information available on Wikipedia to query engines.

The Semantic Web in its macro form may perhaps never be able to deliver completely on its promise. This does not, however, preclude the technologies developed around the Semantic Web being applied in a less grandiose manner, in the form of HyperImage projects, which can be considered as (micro) semantic webs in their own rights. Assuming the appropriate access has been granted, a collection of HyperImage projects in a particular installation could be combined to form a larger, but localised semantic web.

Despite the limitations of the Semantic Web, it might still be possible to derive useful information that can be presented to the users of a particular HyperImage project and vetted according to their professional expertise. The HyperImage Virtual Research Environment is essentially a server-client architecture. In other words, a user accesses her project by starting the client on her local machine, but the data itself is all stored remotely on the server. The server can continually query the Semantic Web in the background, attempting to expand the semantic web created in the project by using the information already gathered to generate new queries and thus perhaps suggests new avenues of research to those involved in the project. The next time the user logs into her project, the system may suggest some new pieces of information that might be relevant to the project. There must always be a human in the loop to verify the new information and decide whether it is valuable enough to be added to the project. An example will help clarify the process.

A HyperImage Example

At the heart of a HyperImage project are objects and their views. For example, an object such as a sculpture may possess several views – photographs or drawings from different perspectives – or a particular oil painting may have been documented before and then after restoration. In the simplest case, an object only has one view. A view is thus an image that represents some aspect of the object. A view can contain complex polygons which can form individual layers. These layers can be enriched with metadata and also function as link anchors.

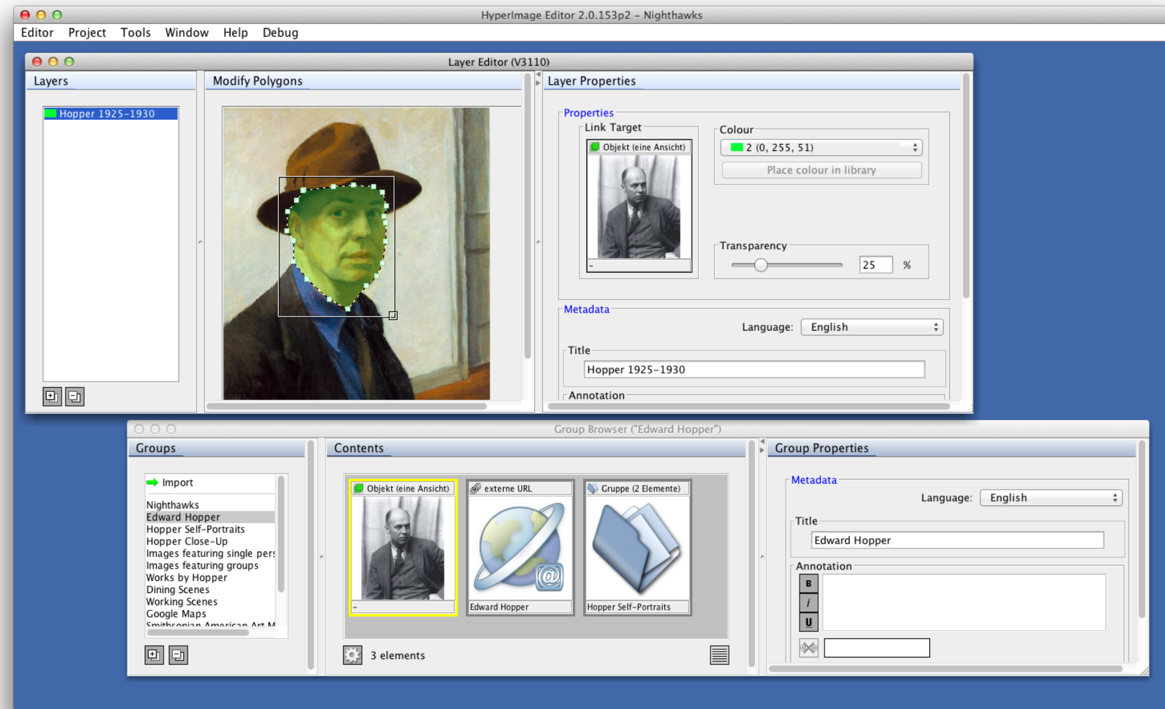

Figure 2: Screenshot of the HyperImage 2.0 Layer Editor

The current version of HyperImage (2.0) supports simple links between anchor and target (or multiple targets if the link points to a group element), but the next version under development will support typed links, thus supporting more sophisticated search functionality.

Elements of a HyperImage project can link to an external website. The Hopper Project might contain photos of the diner considered to have been Hopper's inspiration, using its geographical coordinates to link to Google Maps and thus allowing researchers or other interested parties to view the location in its current form.

In addition to being able to navigate the hyperlinks in a HyperImage project, it is also possible to arrange any and all elements in a project on a light table. This is a form of 'free' association, allowing researchers to juxtapose elements of interest and provide another form of intellectual access to the project. The presentation can also be enriched with textual annotations containing further information as well as additional hyperlinks to provide further context.

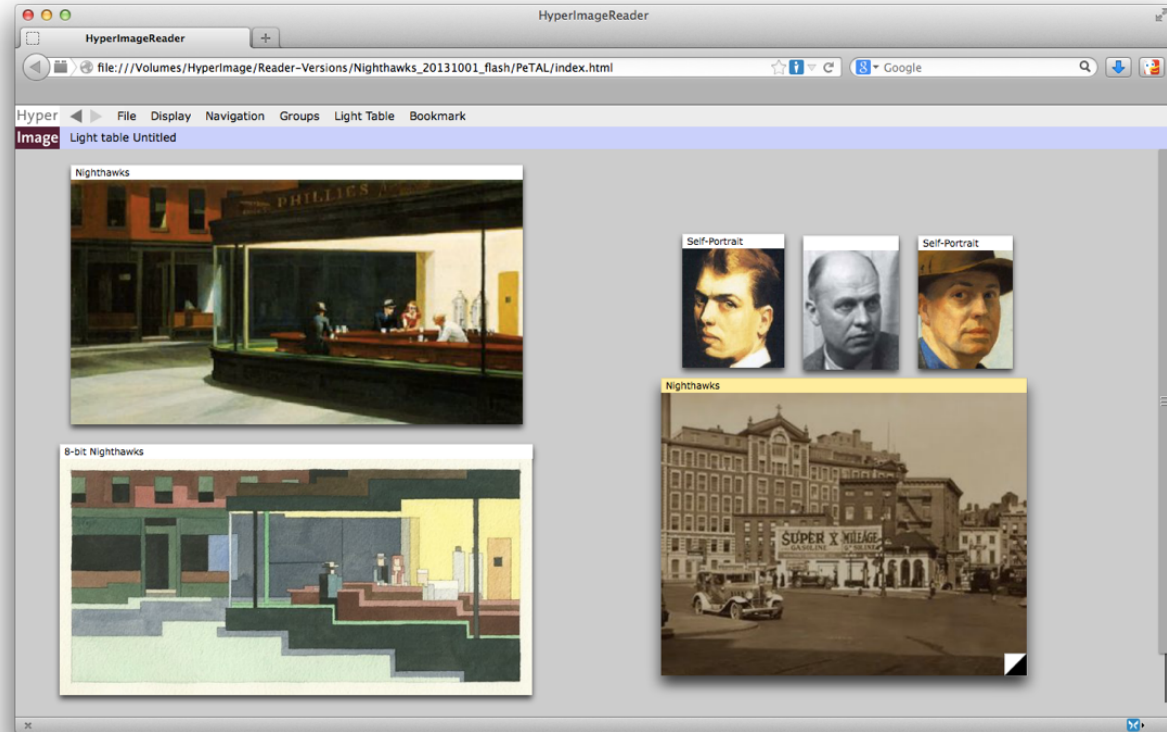

Figure 3: Screenshot of the HyperImage 2.0 Light Table

As the researcher continues to add content to the project, uploading more images, updating the metadata and creating further references and links, the HyperImage server could query various datasets continually in the background. The queries would be based on the content already in the project and the results returned used to suggest new associations that the researcher may not already have made. These results would be presented to the researcher in a discovery mode, allowing data that has been gathered automatically to be reviewed and possibly added to the project. Through positive reinforcement the quality of the data being returned by the automatic queries could be further improved.

The automatic background queries are purely an auxiliary function, augmenting the prime research being undertaken in the project. However, instead of having to initiate searches explicitly, the micro semantic web formed by the project could be used as a basis to make queries against the macro Semantic Web. The more metadata added to a project, the more refined the queries, which results in continually improving results. This iterative refinement can continue until the researcher is satisfied with the project and deems it suitable for release to a wider audience.

HyperImage 3.0

The Semantic Web has its issues, but its technologies can be applied within a HyperImage project or across projects, supporting far more sophisticated search options – types and relations as opposed to just keywords and dates. The current version of HyperImage uses Semantic Web technologies only in a rudimentary fashion. The advanced functionality described above will be part of the feature set of HyperImage 3.0, which is currently under development. See the HyperImage website http://hyperimage.ws/ for more details regarding the roadmap and scheduled releases.

Standards, Sustainability, Long-Term Preservation and Open Source

Research involves a huge expenditure of time and energy, so it is not unreasonable for a researcher to require that the work they have undertaken will be available at least to them and hopefully to a broader public for an indefinite amount of time. Digital research systems represent a particular danger in this regard, especially when proprietary formats are involved, since there is no guarantee that research conducted using a particular software system will be accessible at some point in the future.

The HyperImage software is open source and based on open standards (such as Scalable Vector Graphics, the eXtensible Markup **L**anguage (XML) and the Resource Description Framework). The software also provides a number of default metadata templates based on standards such as Dublin Core, the Getty Categories for the Description of Works of Art and VRA Core. It is also possible to define custom metadata templates. While making a software open source does not necessary ensure its existence – an appropriate level of technical knowledge would be required to install a system on a future computer, assuming the project was no longer being actively updated – it does mean that it is possible to peek under the hood and modify or re-engineer the system so that it could run on future computers. This is not an option for proprietary, closed source software systems.

HyperImage was initially funded as a project by the German Federal Ministry of Education and Research. As is the case with many publicly funded projects, after the funding had ceased the future of the software seemed bleak. We as the maintainers of the project have taken the step to form a company in the hopes of further developing this and other software systems in a supranational fashion and to provide services to a diverse range of tertiary and cultural institutions.

We trust that this robust strategy will thus guarantee the software's longevity and ensure that research conducted using HyperImage will be accessible to future generations.

Bibliography

Berners-Lee, Tim. 1989. *Information Management: A Proposal*.
http://www.w3.org/History/1989/proposal.html

Berners-Lee, Tim. 2007. *Giant Global Graph*. http://dig.csail.mit.edu/breadcrumbs/node/215

Berners-Lee, Tim; James Hendler; Ora Lassila. 2001. *The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities*. In: Scientific American, Vol. 284, No. 5, pp. 34–43.

Butterfield, Herbert. 1957 (first pub. 1949). *The Origins of Modern Science, 1300–1800*. New York: The Macmillan Company.

Kuhn, Thomas S. 1970. *The Structure of Scientific Revolutions*. University of Chicago Press.

Schnapp, Jeffrey; Todd Presner, Peter Lunenfeld et al. 2009. *The Digital Humanities Manifesto 2.0*. http://humanitiesblast.com/manifesto/Manifesto_V2.pdf

W3C. 2002. *Resource Description Framework (RDF): Concepts and Abstract Data Model*. http://www.w3.org/TR/2002/WD-rdf-concepts-20020829/

W3C. 2004. *RDF Primer*. http://www.w3.org/TR/2004/REC-rdf-primer-20040210/

Authors' Note

The authors' contact details and more information regarding the HyperImage Virtual Research Environment, including example projects, can be found on the software's website, http://hyperimage.ws/. A demonstration of the software can be requested via the contact form on the website.