CLAROS – bringing classical art to a global public

Donna Kurtz
Beazley Archive
Oxford University
donna.kurtz@beazley.ox.ac.uk

Greg Parker
Beazley Archive
Oxford University
gregory.parker@beazley.ox.ac.uk

David Shotton
Department of Zoology
Oxford University
david.shotton@zoo.ox.ac.uk

Graham Klyne
Department of Zoology
Oxford University
graham.klyne@zoo.ox.ac.uk

Florian Schroff
Engineering Science
Oxford University
schroff@robots.ox.ac.uk

Andrew Zisserman
Engineering Science
Oxford University
az@robots.ox.ac.uk

Yorick Wilks
Oxford Internet Institute
Oxford University
yorick.wilks@oii.ox.ac.uk

Abstract - CLAROS (Classical Art Research Online Services; www.clarosweb.org) is an international interdisciplinary research initiative led by the University of Oxford (Humanities and Mathematics and Physical Sciences), hosted by the Oxford e-Research Centre (OeRC, www.oerc.ox.ac.uk), and inspired by the Beazley Archive (www.beazley.ox.ac.uk) participating in EU R&D projects. During 2009, a pump-priming grant from the University's Fell Fund enabled CLAROS to integrate on line more than two million records and images held in research centres in Oxford, Paris, Cologne and Berlin. CLAROS uses CIDOC CRM (http://cidoc.ics.forth.gr/), developed under UNESCO's ICOM (http://icom.museum/), to map across datasets and a portfolio of Open Source software to deliver them swiftly to a broad range of global users. Data web applications for integration are being developed by Zoology (http://ibrg.zoo.ox.ac.uk/), image recognition by Engineering Science (www.robots.ox.ac.uk/~vgg/), and artificial intelligence by the Oxford Internet Institute (www.oii.ox.ac.uk). CLAROS will welcome new institutional members and engage with the public to document art and disseminate results.

I. THE CLAROS VISION

The scholarship of classical antiquity pre-dates the Renaissance. The art of ancient Greece and Rome has been collected and displayed in the world's great museums for centuries. Even before the web, it enjoyed global acclaim. The web's potential for broadening access was embraced by the Beazley Archive in the mid-1990s when it participated with some of Europe's greatest museums in the EU project RAMA (www.beazley.ox.ac.uk/archive/rama.htm). Then housed in the University's Ashmolean Museum, the Archive had been documenting classical art in worldwide collections since 1979; unlike the great museums, it already had highly structured datasets and employed a university ICT infrastructure.

Over the past thirty years the Archive and its colleagues – Lexicon Iconographicum Mythologiae Classicae (http://www.mae.u-paris10.fr/limc-france/), Research Sculpture Archive in the University of Cologne (www.arachne.uni-koeln.de), and the German

Archaeological Institute (www.dainst.org/) have created about 20 terabytes of data and images relating to many types of ancient Greek and Roman art and its collection and 'reception' since antiquity. In addition to painted pottery, engraved gems and cameos, sculpture and other ancient media, 'modern' media such as antiquarian books and photographs take the user from antiquity to the present. Since this art was made by literate societies and can bear inscriptions in Greek and Latin, the *Lexicon of Greek Personal Names* (www.lgpn.ox.ac.uk/) also participates in CLAROS, firmly placing art in its ancient cultural context and providing a natural bridge to the large and well-developed epidoc community. (http://epidoc.sourceforge.net/).

In 2000 the Beazley Archive invited these colleagues to begin discussions on how best to integrate their combined digital assets for the public benefit. Over nearly a decade, while the group worked steadily towards its goal without dedicated funding, CIDOC CRM was being developed for the international museum community under the auspices of UNESCO's ICOM, and the concept of a Semantic Web was emerging. A pump-priming grant from the University of Oxford's Fell Fund for data web development enabled the CLAROS Vision to come into focus.

The inter-disciplinary collaboration between the Oxford University Humanities Division and Mathematics and Physical Sciences Division was facilitated by OeRC, the University of Oxford's e-Research Centre, dedicated to bringing academics across the university together. Both Donna Kurtz (Faculty of Classics) and David Shotton (Department of Zoology) are Research Affiliates of OeRC. Its Director, Anne Trefethen, and the University's Director of IT, Paul Jeffreys, have guided and actively encouraged the initiative.

This paper describes the work carried out during 2008/9 and the applications planned for the coming year.

CLAROS has many goals. Among them are: making digital facsimiles of rare and expensive primary reference books available to anyone, anywhere, at any time, enabling scholars to access datasets of intellectually coherent material easily and swiftly through one multi-lingual search facility, enabling museums to access scholarly records about their own and other museums' objects, and permitting the public and educational institutions (schools, colleges, universities) to engage with 'high art' in innovative ways such as image recognition and the creation of virtual exhibitions. Within the art world, CLAROS has the potential to form a global Open Source registry of objects. Beyond it, CLAROS offers the potential of an intellectually sophisticated search schema that can be adapted for Science as well as Art. We hope its model of interdisciplinary collaboration will inspire others.

New institutional members with extensive digital datasets to share will be welcomed within CLAROS; the service is theoretically infinitely extensible. Each member retains his own assets in his own format, his own website and his own IPR. A Collaboration Agreement drafted by the Legal Services of the University of Oxford is offered to institutional partners; they are encouraged to obtain their administration's approval of this document, and to share their comments and suggestions with the CLAROS partnership.

Since classical art is appreciated by a global community, future development of the service will offer the public ways to engage with experts, and also to contribute information about their own privately held data and images. For example, we envision digital images taken by cameras and mobile phones 'on site' uploaded to CLAROS 'clouds' for image recognition, identification and expert documentation. We also envision personal downloads through mobile phones to provide expert documentation on museum objects being viewed, and an intelligent 'Companion' to the rich nexus of material that will automatically extract additional relevant data from the web according to the user's needs. (Donna Kurtz)

II. CLAROS DATASETS

There are five founder members of CLAROS, each with extensive datasets on classical antiquity (mainly art) that have been created over the past thirty years. Collectively they have more than two million records and images that will be available through CLAROS. Each had highly structured data about intellectually coherent material – the material culture of ancient Greece and Rome. That made it technically relatively straightforward to bring them together virtually and academically worth doing. A guiding principle of CLAROS is that no partner should need to change the format of his data to join. Each of the founder members uses different databases and front end programs for entering, querying and displaying results through their own websites. Data have been exported from each partner into a common CIDOC CRM format. The diversity of data formats demonstrates clearly how other institutional partners can join CLAROS easily. David Shotton describes the dataweb technologies in section 3.

A. University of Oxford - Beazley Archive

The Beazley Archive, established in 1970, began to use electronic documentation in 1979, and now has databases of ancient (mainly Greek) pottery (150,000 records and 130,000 images), engraved gems and cameos, antique and post-antique (50,000 records and 50,000 images), plaster casts of classical sculpture (900 records and 1,000 images), antiquarian (900 records and 1,000 images) and topographical photographs (classical and non-classical sites).

Data are stored in an extensible database system that was developed by the Beazley Archive to enable searches to be made across its different datasets. This system also allows researchers to define their own data structure easily and provides a web interface for updating, maintaining and presenting the data. The underlying data are stored as XML in an SQL Server database. Advanced automatic indexing allows the use of this data in a flexible relational manner for rapid data entry and searching. An ASP export utility was developed for the SQL Server database to export the data in CIDOC CRM format.

B. University of Cologne – Research Sculpture Archive, Arachne

The Research Sculpture Archive began to use electronic documentation in 1972 in association with the German Archaeological Institute (**C**). It documents classical sculpture and has databases with 250,000 records and 490,000 images. Its photographic archive has about 160,000 negatives and prints that are being digitized. Arachne works closely with the German Archaeological Institute (see below).

Data are stored in an underlying MySQL database with a PHP front end and exported to XML syntax that can describe the peculiarities of the internal data model of Arachne; relations between multiple database tables are correctly described and Extensible Stylesheet Language Transformations (XSLT) can be made, cleaning data and mapping to RDF/XML. Finding a piece of XSLT code transforms the internal structure to RDF/XML while preserving the intended meaning of the data.

C. German Archaeological Institute

The German Archaeological Institute has its central office in Berlin and Institutes in Rome, Madrid, Athens, Cairo, and elsewhere. It undertakes archaeological excavations and scientific research. Its photographic archives, with more than 1,500,000 images, about 110,000 already digitised) have been created over more than one hundred years and are still being enlarged by new campaigns.

D. University of Paris X - Lexicon Iconographicum Mythologiae Classicae

Lexicon Iconographicum Mythologiae Classicae is an international scientific enterprise that began in 1972 with scholars, academies and research centres from thirty-eight nations that works with museums of classical antiquities in many of these and other countries. From 1981 *Lexicon Iconographicum Mythologiae Classicae* has been published in eight text and eight plates volumes (8,400 pages of text - in English, French, German, or Italian, with 32,000 black-and-white photographs grouped in 5,800 plates), and two volumes of Index (1,026 pages). LIMC central editorial office, located in Basel, is linked with LIMC documentation and research centres in Athens, Heidelberg/Würzburg and Paris. The Paris office has been at the forefront of LIMC electronic documentation, focusing on iconographical documentation from France, East European, North African and Near Eastern Countries.

LIMC documents mythological and religious iconography on different types of classical art and has more than 100,000 records and 180,000 images from about 2,000 museums and private collections (Paris). LIMC archives are being maintained and made accessible on-line in eight languages (English, French, German, Italian, Greek, Russian, Spanish and Arabic).

Data are stored in an underlying MySQL database with a Java front end. (**A**)

E. University of Oxford, Lexicon of Greek Personal Names

The Lexicon of Greek Personal Names (LGPN) was established by the British Academy in 1972 to provide the authoritative collection of ancient Greek personal names, drawing on inscriptions, papyri, literature, coins and artefacts, from the 8th century B.C. down to around 600 AD. Since 1975, LGPN has used electronic means to store, manipulate and publish its data. 300,000 records, sharing over 35,000 different names, have so far been published in six volumes. Each record gives the name of the individual, the place of origin (from a gazetteer of over 3000 places), the date (normalized to allow sorting etc), and bibliographical details. Parent/child relationships and status are also recorded if known.

The data are stored in a relational database which is used to drive the typesetting of printed volumes, and other outputs. The project has also defined an archival/interchange format in XML against a subset of the Text Encoding Initiative scheme, which is delivered using an eXist XML database. This allows querying of the data in a wide variety of ways, with results being delivered in various formats (using XSLT transformations). One of these delivery formats is RDF/XML using the CIDOC CRM ontology, which is used to expose the project through CLAROS. (Greg Parker)

III. THE CLAROS DATA WEB

A. Background

The Image Bioinformatics Research Group in the Zoology Department of Oxford University has been working to apply the technologies and principles of the Web to integrate biological information. This has led to the concept of ***data webs***, which provide subject-oriented integrated views over existing distributed sources of heterogeneous data. A data web is deployed as a lightweight composition of available Web software elements, drawing particularly on the emerging community of ideas, standards and tools for the Semantic Web and the Web of Linked Data. Recently, available software tools have evolved to a state of maturity that allows us to create robust and scalable applications that combine information from independently deployed bioinformatics sources, an example of which can be seen at OpenFlyData.org. Through undertaking this work we have gained a detailed understanding of the problems and opportunities presented by such a method of data integration.

B. Data web principles

The fundamental principle behind such a data web is that each independent resource is presented as a SPARQL endpoint against which queries using SPARQL, the Web query language and protocol, can return information in the RDF format (Resource Description Framework, the Web standard for data representation). RDF provides a syntactic framework that allows automatic data integration. Two problems remain to be addressed. The first is that of semantic integration, requiring alignment of the data schemas employed by the different resources, ideally by mapping to a common schema or ontology, thus ensuring that labels (e.g. database schema class names) meaning the same thing (e.g. 'artist' in one database, and 'creator' in another) are recognized as such. The second is that of co-reference resolution, ensuring that different references to the same thing (e.g. different database references to the same entity or physical object, recorded using different names) are also recognized as such.

C. CLAROSWeb

We have successfully applied our biological data web experience to the creation of a prototype CLAROSWeb, the CLAROS data web, initially to provide cross-database search capability over five existing databases: the Beazley Archive (BA), the Lexicon of Greek Personal Names (LGPN), the central object database of the German Archaeological Institute (DAI), Arachne and the Lexicon Iconographicum Mythologiae Classicae database (Paris).

The architecture of ClarosWeb is such that additional classical art data providers can be integrated into the system at any time, thereby enriching the total content, simply by mapping their schemas to the core ontology and by making appropriate entries in the co-reference service. The main functional components of the CLAROS data web are shown in the figure below.

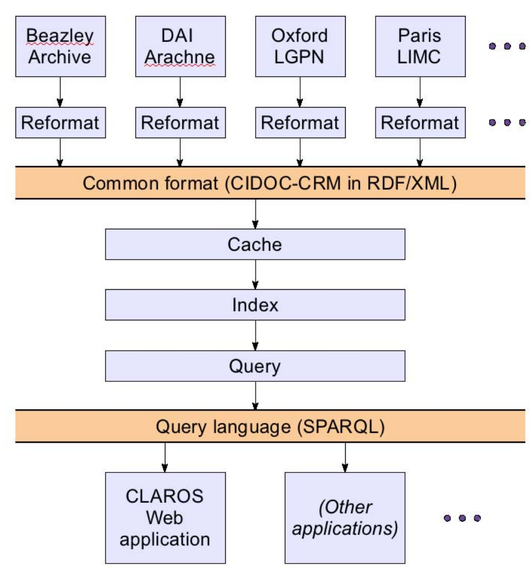

Figure 1. The CLAROS data web, delivering information about classical art objects to on-line users.

- *The schema alignment service* provides information used to map schema terms (i.e. class data) between the source-specific data schemas and the core data web schema.
- *The co-reference service* records alternative names or identifiers for common entities (i.e. instance data) used in different data sources. For example,

names are stored differently by Beazley and LGPN: Beazley stores inscriptions in an Anglicized version of the Greek e.g. Polygnotus, Herakles, while LGPN stores them in Unicode Greek, e.g. Πολύγνωτος, Ἡρακλῆς, and in a strict transliteration including accents.

- *The query handling service* distributes SPARQL queries to each relevant data source, and then integrates the returned information for display to the user.
- *The search and browse services* handle the user interfaces that permit user queries and data browsing.
- The *LIMC multilingual thesaurus* provides relevant terms in English, French, German, Italian, Spanish, Arabic, Russian and Greek, to enable users to query ClarosWeb.

D. Information modelling

A goal of CLAROS is to create a coherent and high fidelity view over its varied data sources. To this end, we needed to find points of common reference between the data sources, for which purpose simple word matching is inadequate. We also needed to know the context of word-use (e.g. as a personal name, as a place name, or describing some activity). The data sources with which we work already contain much of the desired contextual information, but they generally represent it in very different ways, so we needed a common conceptual framework for the confluence of this information. The CIDOC Conceptual Reference Model (CRM) is such a framework, designed for cultural heritage information, which joins philosophical thinking with data management experience. It is also amenable to representation using RDF. We have adopted CIDOC-CRM for CLAROSWeb, since it is well suited for our purposes, and allows common data elements to be used to represent common concepts, without constraining the overall range of information that can be represented.

Using CIDOC-CRM, we have created a common representation for key concepts from the five initial databases, and a basis for including additional data sources in the future. This will, in turn, form the foundation for a service that indexes key concepts from these sources to provide a query facility that operates uniformly across all these data. CIDOC-CRM also supports the targeted use of thesauri for multi-language support, both in the diversity of data sources and in user queries.

E. Data providers

The CLAROS data web will accept and index arbitrary CIDOC-CRM data, so participating data providers need only to present their data using CIDOC-CRM concepts, formatted as RDF. CIDOC-CRM allows information to be presented with varying levels of detail and formality, so our only requirement is that the data providers enable a simple facility that extracts and reformats their data for use with CLAROS. The hardest part is deciding how to represent the original data structures as CIDOC-CRM concepts. However, for this much of the leg work has already been done in creating the initial prototype CLAROS data web.

In practice, new data providers will be free to choose how much (or how little) of their data they expose in this way. CLAROS will index what is provided, and query responses will provide links back to the provider's web sites as appropriate. Clearly, the more metadata that are provided, the more accurate and detailed the possible query responses will be. (David Shotton)

IV. Image Recognition

A. Introduction

Since CLAROS documents objects we want to be able to use an image to query a large corpus of images and return images that contain the query object. In much the same way as Google uses a text query to query a large corpus of text documents (the web) we want results immediately and accurately (meaning that the first page of results at least should be useful and correct). There are two key elements required to achieve this. First, as much processing as possible should be carried out in advance (off-line), so that when the user sends the query all that is necessary to answer the query can be performed very efficiently. In the case of Google an index is constructed in advance where a list is built of all the web pages containing each particular word, e.g. CLAROS. Then at the query stage only the web pages in that list need to be considered further (rather than every web page). The second key element required is a ranking function so that the correct results are returned first.

For querying images with an image, it is not possible to use words directly of course, and instead a descriptor must be determined for each image. The same two elements apply though: the descriptors for all the images in the corpus are computed in advance (so it does not matter too much if this computation is quite costly); and at query time the descriptor must enable accurate and efficient matching of the images. The particular descriptor will vary depending on the object of interest. For example, if the goal is to return vases of the same shape (as will be illustrated here) then the descriptor depends on the silhouette of the vase, but not the decorations. On the other hand if the goal is to return all depictions of a particular god, e.g. Zeus seated, then the descriptor will only depend on the decorations and can largely ignore the shape.

In the sequel we describe how these ideas can be applied to build a real time retrieval engine for Greek vases in the CLAROS collection, and then an application to classifying user uploaded photographs.

B. Describing and matching vases by their shape

Vases are often photographed from the side against a neutral background. In order to represent the shape of the vase silhouette a number of processing stages need to be carried out: first, the vase foreground region is separated from the background. This is achieved by learning the foreground colour from an area around the centre of the image and the background colour from the periphery, and then carrying out an optimization to obtain a connected region with this foreground colour [1, 2]; second, the symmetry axis of the vase is determined as the mid-line of the silhouette; and third, the shape of the vase is recorded as a vector of the widths as a function of the vectical height in the image.

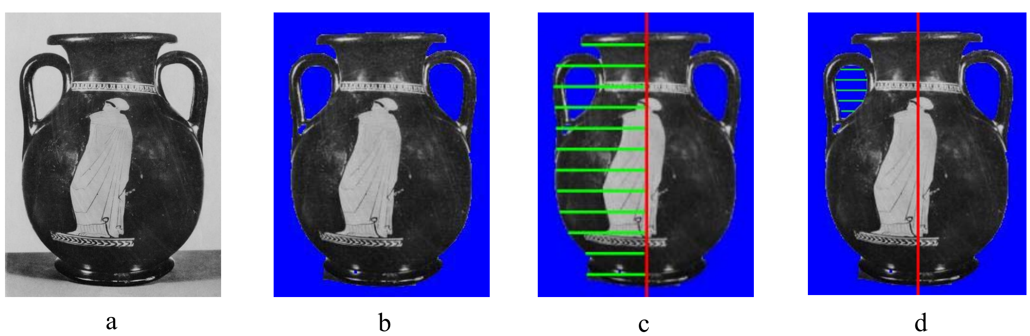

Figure 2. From image to shape descriptor: (a) the original image; (b) separation into foreground (the vase) and background (shown blue); (c) the symmetry axis and a subset of the exterior width measurements; (d) the handle measurements.

Separate measurements are made of each side of the vase and of the interior of the handles. The processing steps are illustrated in figure 2. The outcome is a vector of a 100 or so measurements that describes the vase.

Matching one vase descriptor to another simply involves measuring the distance between the vectors, using a sum of squared distances in the manner of Pythagoras. At run time, our task is to rank all the vases in the corpus by their distance to the query vase. In the Beazley collection there are more than 100,000 images of vases, so matching against all of them, and then ordering by distance can be a costly task. Consequently, an efficient data structure is employed that can very quickly, but only approximately, determine the distances. We use a random forest of KD trees for this [3].

In summary, by first processing all the vases to obtain their shape descriptors, and then employing efficient nearest neighbour matching at run time, we have constructed a retrieval engine that given a query image of a vase can generate immediately (in less than a second) a list ranked by similarity of shape to all the vases in the corpus.

C. Classifying vases from iphone photos

Given the vase search engine it is then possible to engineer a real time system whereby a person can photograph a vase, have it matched to images in the Beazley archive, and thereby identify the shape of the vase and get full documentation on vases with that shape in the Archive. The process is illustrated in figure 3. An image can be uploaded (from a disk or over the internet), its shape computed and classified, and also a ranked list of images from the archive generated (see figure 4). The entire process takes a matter of seconds.

Clicking on any of the matched images gives information on where it was found, its decorations etc. The shape of the uploaded vase is identified by comparing it to the nearest (in descriptor distance) three vases in the archive and, if they agree on shape, then attributing this shape to the query vase. (Florian Schroff and Andrew Zisserman)

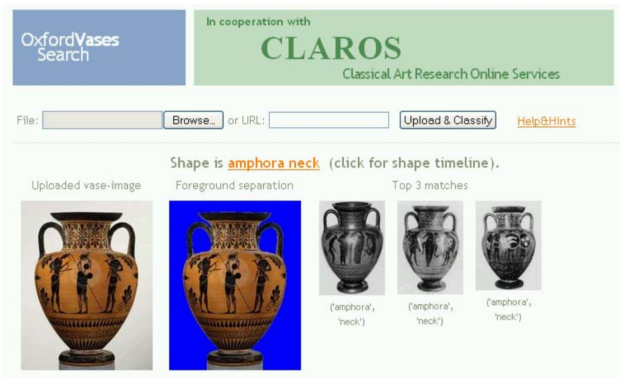

Figure 3. Identifying vases from photos: a screendump of the working system – a vase is uploaded from a web link and its shape classified. The uploaded image is shown on the left, and its shape identified as `amphora neck' (top). Three of the matching vases from the Beazley collection are used to determine the shape (shown on the right).

Figure 4. Matches to the uploaded image in the Beazley archive. The matches are ordered (ranked) in reading order.

V. COMPANIONS

The COMPANIONS (www.companions-project.org) is an EU funded-project involving sixteen partners that aims to change the way we think about the relationships of people to computers and the Internet by developing a virtual conversational 'Companion'. This will be an agent or 'presence' that could stay with the user for long periods of time, developing a relationship and 'knowing' its owners preferences and wishes. It communicates with the user primarily by using and understanding speech, but also using other technologies such as touch screens and sensors; it can in principle take any form, from a mobile phone or PC screen to a cuddly toy or handbag. The website allows you to see some of the trial Companions we have created in the project's first two years: the Senior Companion [4] discusses a person's photograph images and helps organize them, while trying, through informal conversation, to build up a representation of the person's life, relatives, friends and contacts, as well as the places they have visited and their family events.

The question we are exploring is whether a Companion of this sort, communicating about images through conversation, could have a role in organizing information about the classical world and helping disseminate it to a wider audience. Already colleagues at the Oxford Internet Institute (www.oii.ox.ac.uk) are exploring through funded research [5] the application of a Companion in a pedagogical context, managing an individual's learning and creating a model of what the learner knows at any stage.

But our setting is not simply pedagogical; we could think of two different aspects of a Companion faced with a collection of classical art images: one can envisage a Visitor's Companion—rather like a far more intelligent form of the audio tapes one takes round an art gallery and which "knows" which picture one is looking at. A classical Visitor's Companion would explain images, answering questions about background information concerning classical persons and places, possibly making judgments about how well informed a particular user was and tailoring background material so as not to bore him with the obvious. The "talking head" of the Visitor's Companion might be a moving speaking image of a well known classical person: Virgil, of course, was chosen by Dante as his Visitor's Companion through the Inferno.

But one can also envisage a Scholar's Companion, one now closer to the original in that it could stay with the scholar for long periods of time and know where they had been together in a collection. That Companion would not need to offer much background and would not need a famous face: but it would be able to run technical comparisons, say, of a new amphora image against wide sets of amphorae so as to aid the classification of a new image. An interesting possibility here is that the two Companions need not be totally distinct, and one can see the need for a hybrid between them; for example an undergraduate student of classical art and history is somewhere between a scholar and a visitor. Such a student might need to ask questions about the identity of classical persons and dates that would be obvious to a scholar.

The underlying Companions technology is important: some of it is simply specific to the understanding and generation of reasonable speech and the language technologies needed to understand and assimilate what is said to it. All this has advanced considerably in the last twenty years, though it is by no means perfect; a Companion still has to be able to recover when it has misunderstood something, just as a slightly confused person does from time to time.

The more general technology that is of interest is the ability to access not only large databases of images in the collection, but to access in real-time other open sources of knowledge about the classical world such as the Internet, and the ability to build up structures of new relationships that the Companion has helped discover with the scholar/user; these are what are now usually called ontologies, or knowledge structures. The crucial point here is that this kind of reasoning and knowledge building, which was always part of the programme of classical Artificial Intelligence [6] for fifty years, has taken a completely new turn since the development of the Internet and, under names like "The Semantic Web" [7], has become a new paradigm for thinking about and manipulating public knowledge. A key difference is that AI was traditionally limited to the knowledge you had coded explicitly into a system, but now a Companion can reach out to the Internet and seek knowledge on anything, way beyond what the system designer had thought of. Told that a vase comes from Thrace, a Companion can reach out instantly on the Internet to places and dates in the history of Thrace.

The public now encounter this kind of underlying structure every day in applications like Facebook where their own images and news and relationships are organized by a system in an increasingly sophisticated way. The original Companion, talking about a user's pictures, goes off to Facebook entries for the user and their friends to gather more information about them without being asked to do so. A possibility arises that this might, in reverse as it were, be a tool for a Visitor's or Undergraduate's Companion to store electronically what had been discovered during the route through a collection: something they could come back to at any time on their iPhones, say. It could take the form of a "classical Facebook" of images and persons encountered in the collection, with their associated tags making their relationships and ties to classical places and times explicit. A classical Companion could integrate directly to Facebook (an open platform for development) as a dissemination tool for classical knowledge and experience that could be personalized and in a way virtually all young people are now familiar with. (Yorick Wilks)

VI. References

[1] Y. Y. Boykov and M. P. Jolly. "Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images". *Proceedings of the 8th International Conference on Computer Vision*, volume 2, Vancouver, Canada, 2001, pp. 105-112.

[2] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: interactive foreground extraction using iterated graph cuts. *Proceedings of the ACM SIGGRAPH Conference on Computer Graphics*, 23(3), 2004, pp. 309-314.

[3] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. "Object retrieval with large vocabularies and fast spatial matching". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

[4] D. Field, R. Catizone, W. Cheng, A. Dingli, S. Worgan, L. Ye, and Y. Wilks, "The Senior Companion: a Semantic Web Dialogue System, Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems", *AAMAS*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, Budapest, Hungary, 2009.

[5] C. Davies, and R. Eynon, "A Companion for learning in everyday life". In Y. Wilks, (ed.) *Close Engagements with Artificial Companions: Key social, psychological, ethical and design issues*. John Benjamins: Amsterdam, 2009

[6] Y. Wilks, article "Artificial Intelligence" in the *Concise Encylopedia of Pragmatics* (ed.) J. Mey, Pergamon, Amsterdam. 1998.

[7] Y. Wilks, "The Semantic Web: Apotheosis of Annotation, but What Are Its Semantics?" *IEEE Intelligent Systems*, 23, (3), 2008, pp. 41-49.