

Cross-media document annotation and enrichment

Ajay Chakravarthy
University of Sheffield

Regent Court, 211 Portobello Street,
Sheffield, S1 4P, United Kingdom
+44-114-2221945

a.chakravarthy@dcs.shef.ac.uk
k

Fabio Ciravegna
University of Sheffield

Regent Court, 211 Portobello Street,
Sheffield, S1 4P, United Kingdom
+44-114-2221945

f.ciravegna@dcs.shef.ac.uk

Vitaveska Lanfranchi
University of Sheffield

Regent Court, 211 Portobello Street,
Sheffield, S1 4P, United Kingdom
+44-114-2221945

v.lanfranchi@dcs.shef.ac.uk

ABSTRACT

Annotation of documents is a complex and labour intensive task. So far, research has focused on supporting the annotation of documents in single media, e.g. texts or images. Much less attention has been paid to the issue of annotating documents across media, especially useful for web documents that usually contain both text and images.

In this paper we describe AKTiveMedia, a tool which supports human-centric annotation of documents across media. It offers a number of features to support different types of annotations, from ontology-based ones to free comments. We discuss what we believe are the main requirements for annotating Web documents, from support of annotator communities, to the reduction of the annotation burden, to the support of document lifecycle and how they have been implemented inside AKTiveMedia.

The tool has applications in annotation of web pages, personal memories and knowledge management.

Categories and Subject Descriptors

H.5 [Information Systems] Information Interfaces And Presentation

General Terms

Design, Human Factors.

Keywords

Semantic Web Annotations, Image Annotation, Text Annotation, Knowledge Management, Document generation.

1. INTRODUCTION

The amount of multimedia information stored every day in both companies and personal archives is growing. Together with the Web, these archives are reaching sizes unimaginable even until some years ago. For example, in August 2005 Yahoo claimed to cover 20 billion pages¹, of which 19.2 billion web documents, 1.6 billion images, 50 million audio and video files. It is also calculated that almost 375 petabytes or 787.5 billion photographs) are produced each year (almost 2 times all printed material) with a yearly growth rate of 5%, which is the highest

growth rate among different data types [1]. At the same time, large organizations' intranets have reached the size of mini Webs, connecting thousands of computers and having reached dimensions of dozens of millions of documents; it is expected that soon they will reach hundreds of millions of pages, i.e. a size comparable to the Internet at the end of the 90s. Keyword matching based systems are struggling to keep pace with such an amount of information. While currently there is not an issue in indexing and retrieving documents on the Web, in large organizations and in personal archives the issue of efficiently and effectively retrieving material is becoming pressing, due to the density of information.

Also, keyword-based methods are unable to put information in context. This is a problem for knowledge management, where very often it is the context that determines the importance of a document. The context is very difficult to model in keyword based queries. This becomes impossible when part of the context is spread across different media (e.g. in images).

For this reason there is a growing interest in applying methodologies able to capture the content and the context of multimedia documents, in order to enable effective searching (and document-based knowledge management in general).

A common and successful approach to organise and manage huge quantities of information is to enrich documents with metadata. Previous research in personal image management [14, 8] and text annotation [3, 7, 13, 10] demonstrated how annotating images or documents could be a way to organize information and transform it into knowledge that can be used easily later. Metadata enables the creation of a knowledge base which can then be queried as a way both to retrieve documents (via content and context) and to query the structured data (e.g. creating charts illustrating trends).

In this paper we are first of all trying to identify the main requirements for cross-media annotation, introducing then AKTiveMedia, a tool that supports cross-media (image/text) document annotation. We show how AKTiveMedia supports different types of annotations, from ontology-based to free comments, how it supports communities of annotators, and a document lifecycle, allowing users to both create and annotate documents.

The aim of AKTiveMedia is to address one of the main problems of document annotation: the task complexity. In general AKTiveMedia is a tool that fosters knowledge reuse.

¹ <http://google.weblogsinc.com/2005/08/08/yahoo-upgrades-indexes-claims-20-billion-objects/>

.Finally we describe the underlying architecture and draw some conclusions and future work.

2. CROSS-MEDIA ANNOTATION REQUIREMENTS

We identified some initial requirements for cross-media document enrichment. Firstly, we highlight the dimensions of the content that can be enriched with annotations, then we will proceed discussing how to reduce the complexity of the annotation task, and how to support community sharing.

2.1 Annotation levels and types

We identify five main dimensions of information that can be associated to a document via annotation:

1. Resource Metadata, like creation date, time, author, etc.; this type of information is generally provided in a structured form, for example via EXIF data for images, document creation time for texts, HTML metadata for author, etc. It is quite easy to automatically capture metadata and it provides an important knowledge about the context in which the document was produced.
2. Content annotation: which makes content available for retrieval; typically, in literature, content has been represented using ontology-based annotation. This is the most common type of annotation in the semantic web and is generally used to mark-up contingency situations that can change in time. Annotations can be performed across documents and media, i.e. they may relate the text content with part of an image, as mentioned in the examples above.
3. Immutable knowledge about instances (e.g.); this information is generally stored outside the large majority of documents; it will be described in the ontology. Some documents, e.g. descriptive or normative documents, such as dictionaries, etc. can contain immutable knowledge.
4. Informal knowledge about the document or its content. This is generally stored using free text comments that integrate the document content, adding information and knowledge not explicitly mentioned within the document. For example, a user could explain in the comments why a specific formulation was chosen or why a specific hypothesis was pursued; i.e. comments are used to complement the knowledge in the document with knowledge about the process that generated it.
5. Another possible way for annotating documents and images are folksonomies. In our opinion, folksonomies are more interesting for personal use e.g. to annotate pictures to share with friends than for use in knowledge management. In this case the social dimension of sharing is more important than retrieval; there is no need of formal classification of concepts; folksonomies are more a way to attach emotions and memories. In these cases, free annotation (tags and textual descriptions) proves to be more interesting for the users [7], as demonstrated by the success of community-based

image annotation websites as Flickr² or social bookmark managers as De.li.ci.ous³.

2.2 Annotation and document lifecycle

In our opinion, annotation of document should follow the whole document lifecycle, from production to use and be flexible to support the needs of different types of users. In previous research work [3, 2, 14] the annotation task was considered associated mainly to the document production task. However, annotation can happen every time a document is accessed. This is because:

1. The author may want to make available the document content via ontology-based annotation. The author has generally a specific view on the reasons why a document is produced and successively retrieved.
2. The reader may need a different (level of) annotation than the one provided by the author (i.e. to use a different ontology for marking up content or need more details).
3. All users may want to comment on the document in itself or on other comments.

Not all annotations must necessarily be widely available. Some annotations can be personal, others may stay within specific boundaries (e.g. the department or the company), and others can be made available.

2.3 Complexity in annotation

Manually annotating data is a labour intensive and tedious task [2] for a user. It can increase both the time needed for producing a document and the information overload.

Previous literature studies have highlighted the importance of cooperative systems able to ease the annotation process [10] and reduce the information overload. While it is difficult to completely automate the annotation task, because annotations can refer to subjective opinions and memories, it is possible to help users on many fronts, for example automatically extracting metadata from documents using Information Extraction methodologies

2.4 Community Contribution

As experience shows, the importance of user voluntary contributions is fundamental for the creation of a base of knowledge. This makes the difference between success and failure of applications on the Web [12].

We believe that the social perspective is fundamental as it enables the explicitation of implicit knowledge. Such explicitation of knowledge comes very often in informal comments. This information is generally gladly volunteered by both authors and readers (as the experience of GoogleMaps shows, where a gigantic database of information is created by Web users).

² <http://www.flickr.com/>

³ <http://del.icio.us/>

2.5 Ontology Complexity

Ontologies for annotations can be quite complex. Most of the current annotation tools provide a side panel where the ontology is displayed in the form of a tree. Annotation is done by selecting an element from the tree. This is clearly an impossible strategy with a very large ontology, as the user would have to scroll over a very large tree.

Moreover, a large ontology (even in terms of 100s of concepts) is difficult to use because users find difficult to remember all the available concepts and to use them properly. As previous literature proved [9], when dealing with vast quantities of information users may want to zoom and visualise only the sections they are interested in, or filter out what is not relevant for the current task.

For this reason it is important to find a way to represent the ontology making it manageable when annotating documents.

3. AKTIVEMEDIA

AKTive Media is a user centric system for document enrichment across media; it uses Semantic Web and language technologies for acquiring, storing and reusing knowledge. The aim is to provide a seamless interface that guides users through the annotation process, reducing the complexity of their task.

In the following paragraph we will detail how AKTiveMedia answers to the previously outlined requirements.

3.1 Overview

AKTiveMedia supports the annotation of text, images and HTML documents (containing both text and images) using both ontology-based and free-text annotations.

Support is provided both for author and reader annotations, giving the possibility to load different ontologies accordingly to the task. Moreover the annotations are stored separately from the document, alongside with the authorship. This enables controlling the privacy of annotations and the display.

In order to support community sharing, AKTiveMedia allows the user to insert comments and annotations and share them with other members of the community through a centralised server (more details is Section 4).

Human Language technologies have been employed to ease the annotation task: an underlying Information Extraction systems (T-Rex) [11] has been integrated, that learns from previous annotations (both user and community ones) and suggests new annotation to the user, that can accept or reject them, thus re-training T-Rex and improving the learning process. This is a route we already successfully explored in Melita[2], OntoMat[3] and MnM[14], where it was found that the annotation time could decrease by 80% and interannotator agreement could double [6]. But AKTiveMedia goes beyond the single media annotation suggestions and moves towards cross-media strategies. When an annotation is inserted in the text, the system automatically inserts it in a knowledge base that will be used to suggest new annotations when dealing with images (more details in Section 3.1).

Moreover some metadata are automatically captured via the automatic extraction of EXIF data in images and by extracting meta-tags from HTML documents. In one application, we also integrated GPS and calendar information [13].

The way in which AKTiveMedia deals with the problem of the ontology, is by using what we call “disappearing ontology” i.e. we try to hide the unnecessary complexity of the ontology. On the one hand, users can adopt specific views on the ontology to annotate their documents: a user may not need to use the complete ontology all the time.

Therefore a very high level description of the ontology is displayed and the details are hidden till the user needs to use them. Concepts that are not displayed directly in the graph are retrieved using the search mechanism associated to the ontology based annotation. Inputting a textual description (e.g. “Director of KMI”) and selecting from a short list of potential ontology concepts becomes quite easy for a user (e.g. selecting between “Position” and “Person”). This reduces the necessity of displaying a large ontology, while maintaining its details.

AKTiveMedia supports also editing documents or creating new ones: this is particularly relevant when wanting to create a new Semantic Web Website.

In a previous research project, AKTiveDoc [8], we provided facilities for annotating documents while editing. This enabled semi-automatic searching of relevant knowledge to insert into the document. We decided to avoid this feature (annotating while editing) in AKTiveMedia at this point in time, because in the use of AKTiveDoc we found that it was not an easily manageable feature. On the one hand there is an objective complexity in the software implementation (necessity of aligning annotation and document during modifications).

On the other hand, users found it difficult to manage the simultaneous tasks of editing, annotating and retrieval of further knowledge. It implied a high cognitive complexity which was difficult to manage in a single environment. Users found easier to perform the three tasks using three separate environments. This is why AKTiveMedia separates the editing and the annotation task: a user can firstly create a new HTML document containing both text and images, and afterwards the annotation can be performed as for a normal document.

The editing functionality has been implemented integrating an HTML editor based on EKIT⁴ (Figure 1).

⁴<http://www.hexidec.com/ekit.php>

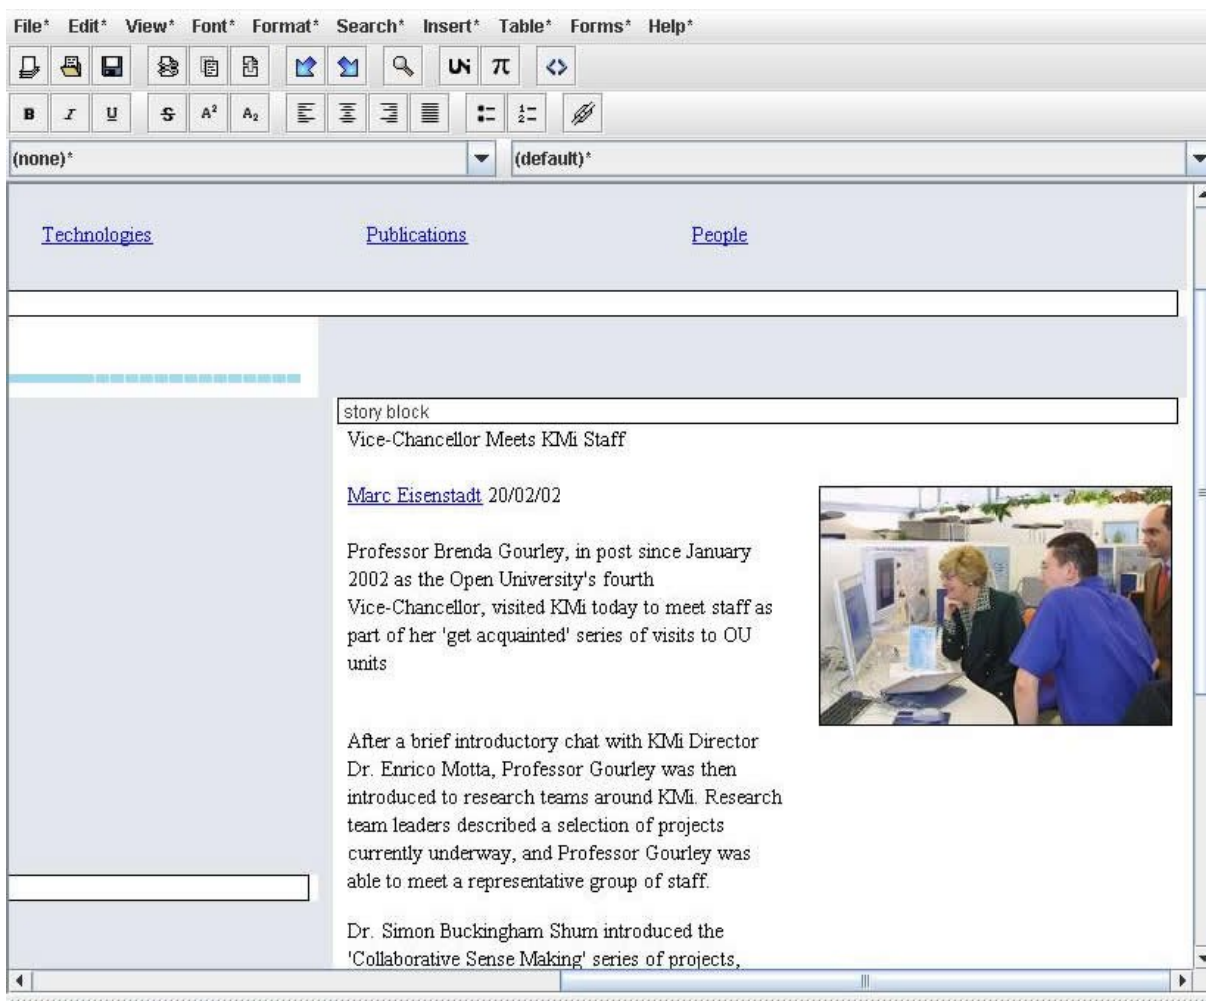


Figure 1 - Editing a multimedia document in AKTiveMedia

The document writing task is also supported via functionalities for retrieving the content of existing documents that operate on the document annotations. This enables reuse of related knowledge. As an example, while the user is writing a document, if he's writing "John Domingue" the system can start retrieving all John Domingue's pictures: the user can then decide to insert one of the picture in the document. This facility enables knowledge reuse making easier for the user to write a document.

Moreover while the user perceives to be writing a simple HTML document, more information is automatically added by the system, in terms of metadata and structure. The resultant file will have an associated RDF annotations file that will contain annotation that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '04, Month 1–2, 2004, City, State, Country.
 Copyright 2004 ACM 1-58113-000-0/00/0004...\$5.00.

the system inserted in a transparent way: for example when creating a H1 heading, the system will match the HTML tag to the document ontology and will automatically insert an annotation "title" with the value inserted by the user.

When the editing is finished, the document can then be annotated as described in Section 3.1.

In the following section we will detail AKTiveMedia interface, using a sample scenario that will allow showing the functionalities. The chosen scenario is the annotation of the KMI news corpus, which is a set of news (published on a website) about the various people visiting the KMI institute and there contributions. The documents are HTML files containing both images and text: content of text and images are related, as they are produced contextually and often describe the same person or event

3.2 Interface

AKTiveMedia supports three main modalities of annotation:

1. images

2. text
3. cross text/images e.g. Annotation of HTML pages.

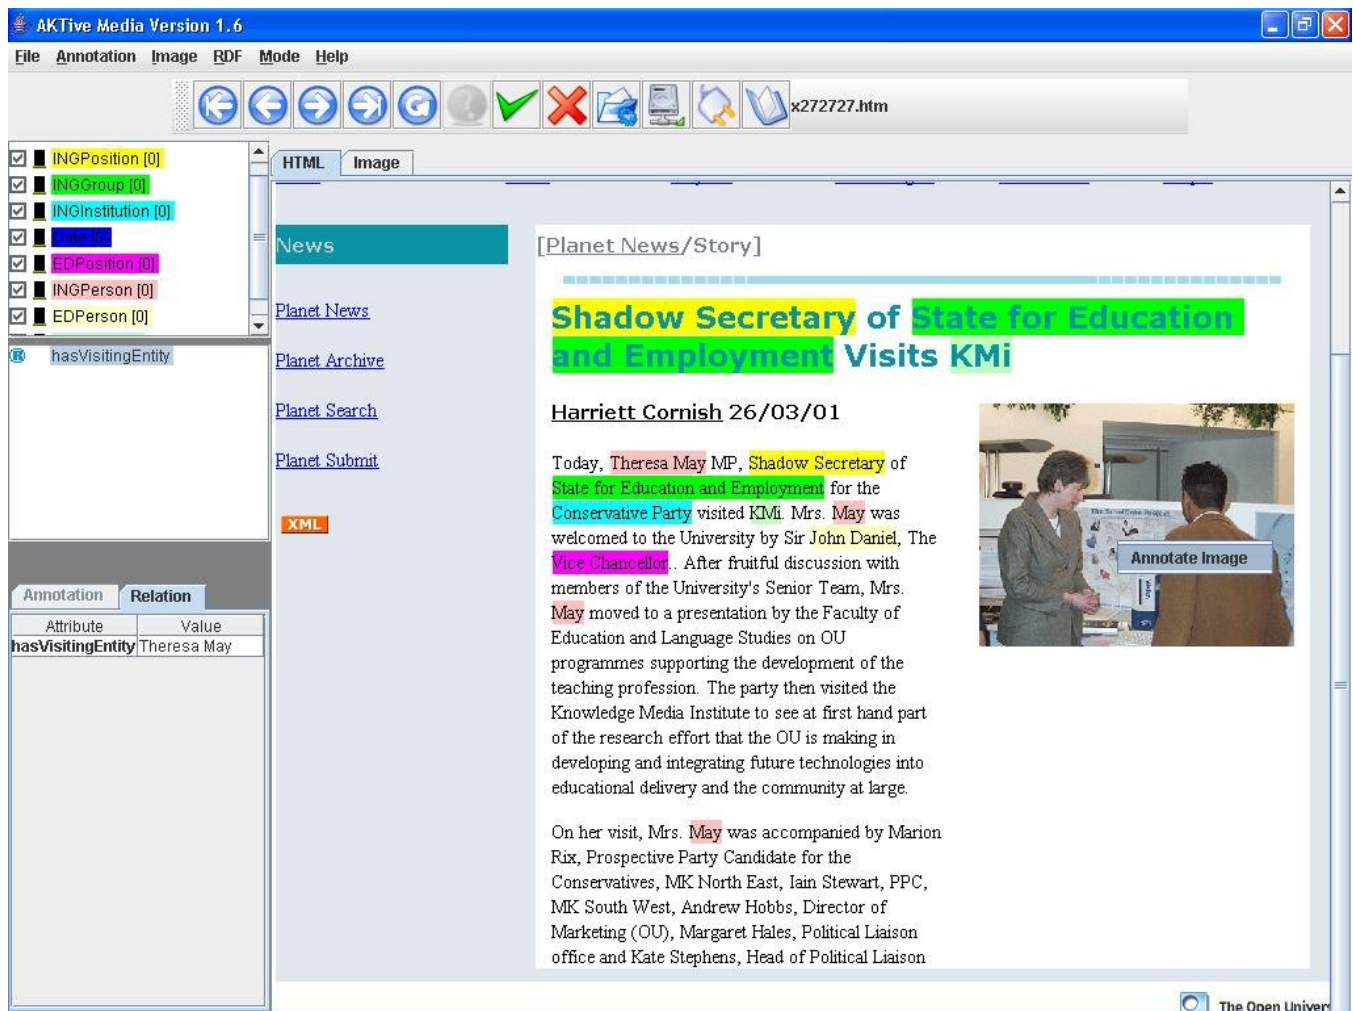


Figure 2 - Annotating relations in HTML document

All modalities share the same functionalities and very similar interfaces. They all offer ontology-based enrichment through a graphical interface, following the paradigm of other annotation tools like Cream or Melita [3, 2]: portions of text or images can be associated with concepts in the ontology with a straightforward point&click interface.

Free-text annotations can also be added on top of the ontology-based ones, to insert more information.

When the user is starting to annotate a new HTML document, they can simply annotate by clicking on the concept in the ontology and then highlighting a sequence of words (see Figure 2).

It is possible to associate relations among highlighted parts (in the figure below we declare again that the “John Domingue” has a visiting entity “Theresa May”). This is done by clicking on the “John Domingue” instance in the text and then on its relation

(“hasVisitingEntity”, mid-left of figure) and then clicking on the instance of the “Theresa May” highlighted in the text (see Figure 2).

When the text has all been annotated, the user can decide to annotate also the corresponding image(s). When right clicking on the image, AKTiveMedia switches to the image annotation mode, without losing the context of the document (the image is opened in another tab).

In AKTiveMedia images can be annotated as a whole or in part. First of all, a title and a description can be inserted for each image, alongside with free text comments related to the whole image. This metadata can be further annotated, using the text based annotation strategy described before. Portions of the image are identified using the mouse (e.g. by drawing a square) and can be annotated via the ontology

When a portion of the image is annotated, a popup window appears (centre right of figure) which enables to describe the content of the annotation in natural language. A facility is provided to search for a complete unique description given the user description. This accesses a triple store of descriptions (as offered by a gazetteer, or by part of the ontology not shown for usability reasons). The facility is used for example to input “John” as the visited person and retrieve complete descriptions of all the persons with the name John working at the KMI institute. The selection has a side effect the allocation of the URI in order to uniquely identify the object.

Ontological relations between instances (i.e. parts of the image) can also be inserted. For example it is possible to declare that the “John Domingue” has a visiting entity “Theresa May”. This is done by clicking on the instance bearing the relation (“John Domingue”); the system will then show all the relations possible for that instance (middle, left in the figure); clicking on the selected relation and then on the other instance (the “Theresa May”) will fill the relation with the latter.

More important, the system allows to establish relations between text and images, allowing the user to assert that the “John

Domingue” in the picture is the same John “Domingue” that was annotated as “VisitedPerson” in the text.

Also these implicit relations can be used by the system to make the annotation task easier, using a contextual annotation mechanisms that analyse user’s or system’s annotations in the text or in the image to suggest new annotations: for example if the text has been annotated first, when the user annotates an area in the image as a *visiting person*, the system will suggest as description the finding previously inserted in the text (“Theresa May”) – (see Figure 3)

The user can accept or reject the suggestion. If accepted, an identity is established between the instances in the text and the ones in the image (same URI).

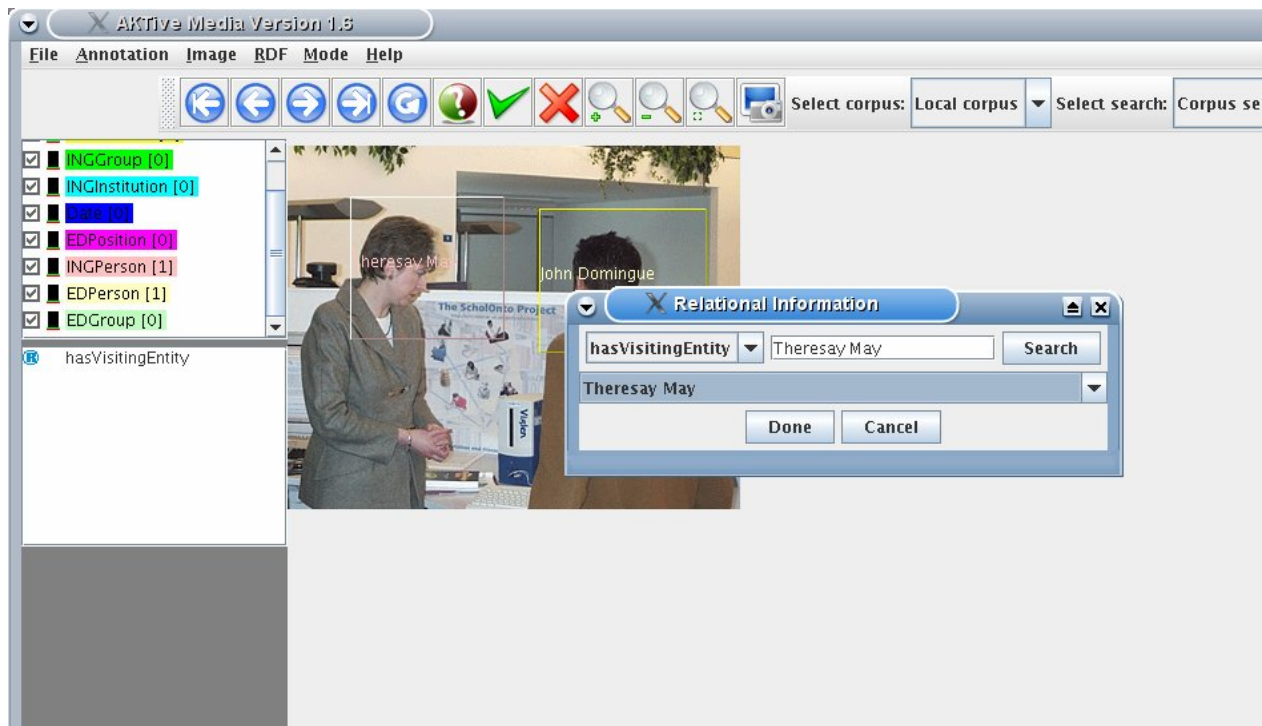


Figure 3 - Annotating an image using contextual suggestions

In case the image is annotated first, the system will search the text for descriptions compatible with those in the image. Matching is done using string distance metrics⁵

4. SYSTEM ARCHITECTURE

The system is based on a configurable plug-in model in which the different components (e.g. ontology loader, annotation modalities,

web services etc.) are independent sub-models that can be plugged in for creating a custom application.

This is because AKTiveMedia is more than just an annotation tool. It is designed to be inserted into user applications. Its architecture focuses on RDF as a way to store and query data and to communicate between components and web services as a way to distribute the architecture. All the annotations are stored as RDF triples inside a local store and periodically updated into a central triple store using web services.

5 <http://www.dcs.shef.ac.uk/~sam/simmetrics.html>

This modular architecture implements the knowledge sharing scenario, allowing different users to see other people's annotations and reuse them (see Figure 4).

When annotations are saved, they are associated to the document through a unique URI or hash code, thus enabling retrieval of annotations performed by other users on the same document.

A plug-in (AKTiveSearch) enables searching and reuse of knowledge while creating or annotating the document. AKTiveSearch enables simultaneous multiple queries to different archives and sources of information, the integration of the returned information and the filtering of the results based upon the context of use.

As mentioned before, AKTiveMedia tries to maximise the amount of resource metadata that can be automatically collected. For this reason an EXIF Extractor and an Information Extraction system, T-Rex, are integrated. They both work in the background, extracting possible metadata and annotations that are later presented to the user and saved in RDF format. In particular T-Rex has been implemented for background training and annotations, using a separate thread, so to not interfere with the user's activity (as the training process can be very long) and to maximise the efficiency. The schema followed is the same as Melita's[2].

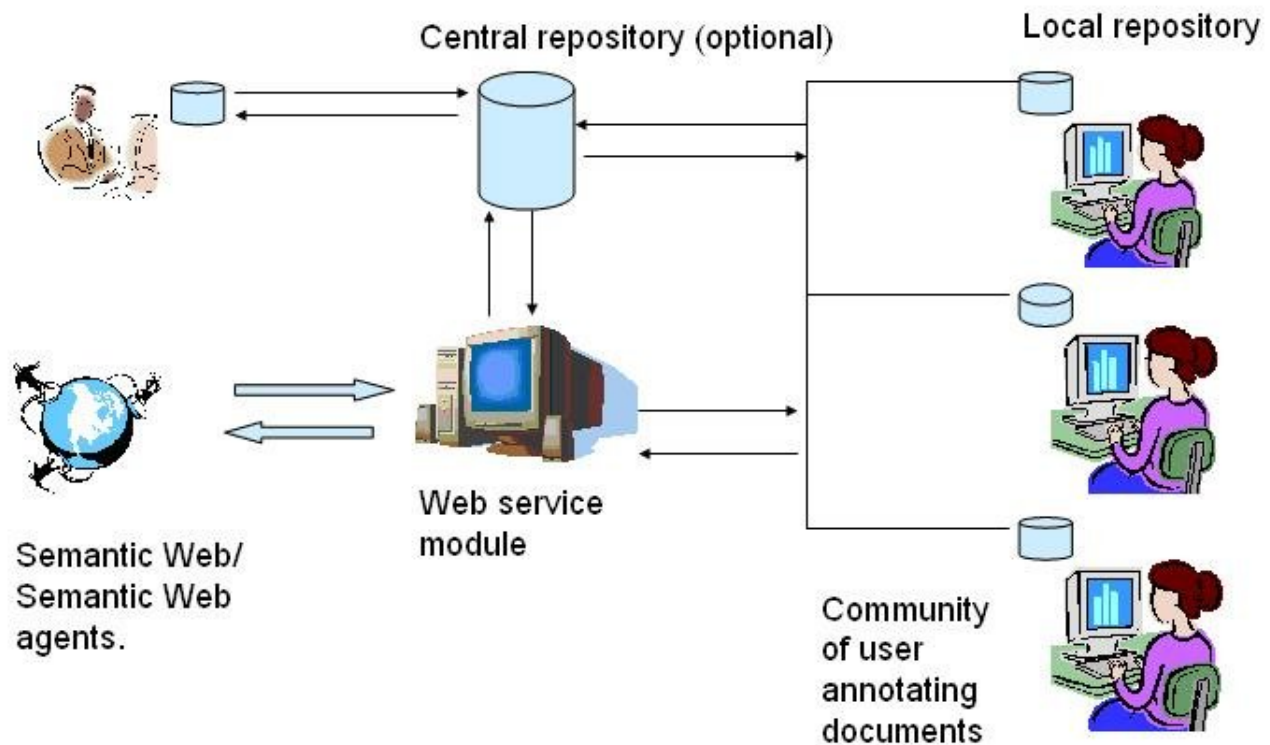


Figure 4 - AKTiveMedia sharing model

The ontology loader, based on Jena⁶, is used to load the users preferred ontology and to implement the selective view mechanism (disappearing ontology). The interaction between the user and the system is realised through the user interface that is also modular to allow different modalities of annotations.

Currently AKTiveMedia supports four annotation modalities (text annotation, image annotation, 3D annotation and Editing) and it is possible to mix and match these modalities in order to

facilitate cross media annotation. The interface component was designed using the MVC (Model View Component) architecture in Java. This enables separation of data and visualization, enabling efficient flow of information across different modalities, while keeping the user interface simple and easy to use.

5. EVALUATION

We have performed a detailed evaluation of AKTive Media during the Fourth Summer School for Ontological Engineering in the Semantic Web in Cerdilla, Spain.

⁶ <http://jena.sourceforge.net/>

Over 60 students divided in groups of 3-4 persons each were testing the system. The task involved first of all annotating 10 documents from the KMI corpus and then starting the Information Extraction system (T-Rex); after the training the system would start suggesting possible annotations in a semi supervised way. Log file were collected, recording all the user activities and students were asked to fill in a questionnaire at the end of the session.

The results of the evaluation are still under study.

6. POTENTIAL USE CASES

In the following sections potential use case in which AKTiveMedia can contribute will be outlined.

6.1 (PhotoCopain) Memories for Life

Memories for Life is a Grand Challenge for Computing Science proposed by the UK Computing Research Committee.

Individuals are usually storing an enormous amount of information about themselves on their computers (documents, images, web browsing logs, etc). The challenge for computing researchers is to develop ideas and techniques that help people get the maximum benefit from their memories, while at the same time giving them complete control over memories so as to preserve their privacy. *Memories for Life* is also regarded as a Grand Challenge by the UK Foresight Cognitive Systems project.

Digital memories clearly offer tremendous potential for science and technology. We must also ensure that they help society by widening access to information technology, so that everyone, not just well-educated people with no disabilities in rich countries, could benefit from the information revolution. The challenge is to develop detailed models of an individual's abilities, skills, and preference by analysing his or her digital memories; and to use these models to optimise computer systems for individuals. A longer-term challenge might be presenting a story extracted from memories in different modalities according to ability and preference; for example, as an oral narrative in the user's native language, or as a purely visual virtual reality reconstruction for people such as aphasics who have problems understanding language. Limited examples of such systems can be built now; the challenge is in mining the wealth of information latent in digital memories so that fully competent systems could be in use in fifteen years.

AKTive Media is being extensively being used for the PhotoCopain project which is a part of the memories for life. Images can be semantically annotated and the narratives linked to the images. This is possible due to the cross media annotation capability of AKTive Media. We are currently focusing on auto narrative generation technology given a set of images in a timeline [12].

6.2 E-Response AKT Project

AKTive Media is currently being ported to integrate with Compendium⁷ tool for the AKT E-Response project. The project aims to use Semantic web agents to automatically deal with an emergency event. Examples of which can be include: taking photographs of the incident and sending them to a semantic web service, locating and notifying the nearest fire stations about the incidents, etc.

AKTive Media will serve as the interface for photographs taken in an emergency situation and there annotations. Further it will also act as a search interface for the photographs using the SPARQL search facility.

7. FUTURE WORK AND CONCLUSIONS

In this paper, we have described and discussed AKTiveMedia, a tool for editing and annotating multimedia document containing images and text. The annotation can be performed within and across the different media. Annotation is mainly manual, but a number of strategies are used to reduce the burden of annotation. We have shown and discussed how the system satisfies a range of user requirements for use to support knowledge management and personal archives. In particular the requirements we analysed are: (1) annotation types and levels, (2) annotation as community activity, (3) annotation and document lifecycle, (4) annotation complexity, (5) ontology complexity and (6) knowledge reuse.

The current applications of AKTiveMedia are in both personal memory management [8] and knowledge management. Concerning the latter, AKTiveMedia is the basis for a real world application under definition within IPAS, a project co-funded by the UK Department of Trade and Industry and Rolls Royce plc (www.3worlds.org). The application concerns the editing and annotation of diagnostic reports on jet engines; the examples used in this paper are derived from the user requirement analysis in IPAS.

In the future, we will explore further levels of community annotations, by addressing in particular issues such as privacy of data and ownership of annotation. Moreover, we will explore in more details the use of folksonomies in an industrial environment, and study their impact on knowledge retrieval and reuse. For this reason we plan to introduce in AKTiveMedia facilities for the direct manipulations of folksonomies. Another venue of development is the annotation of 3D images. The currently available facility implemented in the system is quite limited and needs extensions. In 3D annotations, there is an inherent HCI complexity in annotating an image that can be rotated.

8. ACKNOWLEDGMENTS

This research was partially funded by the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration

⁷ <http://www.aktors.org/technologies/compendium/>

(IRC). AKT is sponsored by EPSRC, grant number GR/N15764/01. This project is also funded by the European project IST X-Media, funded as part of Framework 6, grant number FP6-26978.

AKTiveMedia is an open source project, available under Academic Free License, Educational Community License, General Public License. A copy of the binary files and the source code can be downloaded at <http://sourceforge.net/projects/aktivemedia/>.

9. REFERENCES

- [1] Brilakis, I.: *Content based integration of construction site images in aec/fm model based systems*, PhD thesis, University of Illinois at Urbana-Champaign, 2005, http://www-personal.umich.edu/~brilakis/webdata/Brilakis_Thesis.pdf
- [2] Ciravegna F., Dingli A., Petrelli D. and Wilks Y.: *User-System Cooperation in Document Annotation based on Information Extraction*. In Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), 1-4 October 2002 - Siguenza (Spain), Lecture Notes in Artificial Intelligence 2473, Springer V.
- [3] Handschuh S., Staab S., and Ciravegna F.: *S-CREAM-Semi-automatic CREATION of Metadata*. In Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), 1-4 October 2002 - Siguenza (Spain), Lecture Notes in Artificial Intelligence 2473, Springer Verlag
- [4] Harris, Z.: *Mathematical Structures of Language*. Wiley-Interscience, New York, 1968.
- [5] Iria, J., Ireson, N. and Ciravegna, F.: *An Experimental Study on Boundary Classification Algorithms for Information Extraction using SVM*. Proceedings of the EACL Workshop on Adaptive Text Extraction and Mining (ATEM 2006), April 2006, Trento, Italy.
- [6] Kang, H. and Shneiderman, B. *Visualization Methods for Personal Photo Collections: Browsing and Searching in the PhotoFinder* In Proc. IEEE International Conference on Multimedia and Expo (ICME2000), New York City, New York.
- [7] Kuchinsky A., Pering C., Creech M.L., Freeze D., Serra B., Gwizdka J., *FotoFile: A Consumer Multimedia Organization and Retrieval System*, Proceedings of ACM CHI99 Conference on Human Factors in Computing Systems, 496-503, 1999.
- [8] Lanfranchi V., Ciravegna F., Petrelli D. *Semantic Web-based Document: Editing and Browsing in AktiveDoc*, Proceedings of the 2nd European Semantic Web Conference, Heraklion, Greece, May 29-June 1, 2005
- [9] O'Really, T.: *What is Web 2.0, Design Patterns and Business Models for the Next Generation of Software*.
- [10] Petrelli, D, Lanfranchi, V., Ciravegna, F.: *Working Out a Common Task: Design and Evaluation of User-Intelligent System Collaboration*. In Proceedings of Tenth IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2005), Rome, September 2005.
- [11] Thorsten B. *Inter-annotator agreement for a German Newspaper Corpus*. In the Proceedings of 2nd international conference of Language Resources and Evaluation (LREC 2000).
- [12] *Towards the Narrative Annotation of Personal Information and Gaming Environments* Tuffield, Mr Mischa M and Millard, Dr David E and Shadbolt, Professor Nigel R (2005) *Towards the Narrative Annotation of Personal Information and Gaming Environments*. In Proceedings Hypertext 05, Workshop on Narrative, Musical, Cinematic and Gaming Hyperstructure, Salzburg, Austria
- [13] Tuffield M., Harris S., Duplaw D., Chakravarthy A., Brewster C., Gibbins N., O'Hara K., Ciravegna F., Sleeman D., Snadbolt N., Wilks Y. *Image Annotation with Photocopain* in the Proceedings of the fifteenth world wide web conference (www06), Edinburgh, May 2006.
- [14] Vargas-Vera M., Motta E., Domingue J., Lanzoni M., Stutt A., Ciravegna F.: *MnM: Ontology driven semi-automatic or automatic support for semantic markup*. In Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management, EKAW02. Springer Verlag, 2002