

CULTOS* - Ontological Modelling for Intertextual Studies and the Creation of Multimedia Units of Knowledge

Wernher Behrendt¹, Ziva Ben-Porat², Siegfried Reich¹, Motti Benari²

¹ Salzburg Research, Austria, www.salzburgresearch.at

² Tel Aviv University, Israel, Porter Institute of Semiotics

Contact: wernher.behrendt@salzburgresearch.at

Abstract

We report on lessons learned in developing a multimedia authoring tool which uses an ontology for annotating the media and for defining relationships between segments of the multimedia assets. The hyperlinked knowledge and media structure is stored in an "enhanced multimedia meta object" or EMMO. Adding presentational information can further enhance the knowledge and media structure of the EMMO. A transformer acts as an interpreter of the EMMO structure and creates from it, a hypermedia presentation or textual output. We report on five main lessons learned in CULTOS: lesson one is that advanced end-user software for *non-IT experts* has to satisfy high usability expectations otherwise comparisons with traditional software are just not realistic, because the "advancedness" is offset by its immaturity in usability terms. Lesson two is that Semantic Web applications require ontologies based on community consensus. There is a need for (software supported) platforms that enable such a consensus process to take place. There is also a need to firstly, offer ways to integrate domain ontologies into upper ontologies, and secondly to allow for local (monotonic) extensions to the ontology. Lesson three concerns the quality of the ontologies, in epistemological terms as well as with respect to their formal properties. Future semantic web application development will need to define processes and levels of formality so that one can make rational choices of the trade-offs between formal rigour and practical necessity. Lesson four is that once some knowledge is put into a structure, it remains unclear as to how one can present it in intelligible ways. Transformers for knowledge presentation are in their infancy and we are thinking of ways in which the EMMO is transformed into a media supported hyper-textual narrative. Lesson five is that organisational and infrastructure issues must be taken seriously for any uptake in the commercial field. Our users have created beautiful multimedia showcases (in terms of complexity of knowledge, as well as media assets). The practical limitations lie e.g. in regulatory issues: where can we demonstrate knowledge relationships when the source or target medium is not owned by those who have the knowledge about them, but those who own the copyright? For the infrastructure, the conclusion is that knowledge annotation is by far not enough: in order for a knowledge economy to evolve, a knowledge content value chain (or network) must be supported. Authoring tools support just one phase of that chain. The application domain of CULTOS is the humanities, specifically the field of intertextual research in which semioticians, literature scholars, linguists and cognitive psychologists try to trace the themes, ideas, cultural stereotypes, and medial transformations that have made it possible for cultural canons as well as new paradigms to evolve over time. The project has fostered intensive dialogue between technical groups and a highly motivated user group. Both groups have gone through peaks and troughs as the software improved and expectations grew, neither always fully synchronised with the other. This paper first introduces the vision of CULTOS within cultural heritage and then expands on the lessons learned.

The Vision of CULTOS - Authoring Tools for Preserving and Enhancing a Cultural Canon

Cultural memory is not the sum total of documents in archives, libraries, and museums, although it is often conceptualised in this manner. Just as individual memory is the past as present in someone's mind, so cultural memory is that knowledge which members of any culture share as a human group at any given moment.

No group can survive without it. Any particular language is its most important manifestation. All cultures develop tools for producing an adequate cultural memory. Adequacy implies economy and accessibility. These tools are, therefore, basically mechanisms of selection and reduction: only that which serves a purpose and can easily be remembered (on the basis of a few major attributes repeatedly used) by a heterogeneous population will be retained in cultural memory and propagated through different channels of cultural transmission. Exclusion and reduction are not only instruments of

* CULTOS - Cultural Units of Learning - Tools and Services, IST-2000-28134, www.cultos.org

power (though this cannot be denied) but are also the conditions imposed on cultural memory by cognitive and objective social constraints.

A cultural canon is a particularly important tool for the construction and dissemination of cultural memory.

The CULTOS authoring software empowers a user to annotate multimedia assets and to assign conceptual relationships between the multimedia assets. The annotations are typed with respect to a previously defined ontology. In the case of CULTOS, the ontology describes a hierarchy of intertextual relations.

The CULTOS environment is based around the notion of a specific data structure called EMMO - enhanced multimedia meta object. The data model of EMMOs can be formally described, and offers transformations for different storage layers, as well as defined interface categories that help in standardisation of knowledge assets and their attendant creation, storage and distribution technologies.

A Grand Vision: LEIT, the proposed Library Explicating Intertextual Threads

In Spring 2000, a group of literary experts led by Ziva Ben Porat from Tel-Aviv University proposed the idea of a large digital library consisting not only of digitised works of art, but putting emphasis on the knowledge fabric that links these works of art together, that is, their intertextual relations. Intertextual relations have many similarities with the notion of "memes" (Blackmore 2000) in that they characterise ideas that get taken up and changed, new twists introduced into stereotypical story lines, archetypes used and re-used, culturally coloured re-readings of foreign texts, and relationships between different media, such as the paintings of Velazques based on the story of Cervantes' Don Quixote, or the many rewritings - often as parodies and caricatures- of the windmill scene in Don Quixote. What seems impenetrable to the non expert, is in fact, a fairly well agreed set of known interpretative patterns that differ from group to group in the same way as in IT, database researchers have different approaches depending on whether they work on relational, object-oriented and deductive databases.

The ensuing discussion about LEIT led to the realisation that without productivity tools, such a task was not possible, particularly in view of the fact that there may be many intertextual links between works of art, that have not even been discovered and where one would need very advanced classification mechanisms to detect possible relationships (difficult enough in texts, practically impossible across the arts (how do you detect that a painting is a retelling of a written story, or that a story is inspired by a painting?). At the same time, the community of arts expert was not to be expected to learn HTML to do manual linking and even if that were feasible, then the lack of a common vocabulary would render the exercise pointless. The upshot was the decision to build a prototype multimedia authoring tool that could handle any ontology and any set of multimedia documents, and would allow a community to first define their common vocabulary (at the structural rigour of a computational ontology) and then create multimedia showcases as concrete instantiations of the ontological structure (aka the schema or "data model").

The grand vision of LEIT will need more infrastructure than just the authoring tools of CULTOS: Firstly, ontologies - even when they are stable in principle - will be extended (specialised) in some areas in order to account for specific phenomena. This does not render the old ontology useless, it just adds detail to it. Secondly, there is a need for supporting scholarly dialogue about the intertextual threads once they are submitted. Thirdly, at any one time, people will need "narratives" about threads of knowledge - can machines achieve this? Fourthly and futuristically, as more and more intertextual relationships become structurally understood, art may become a playground for hide and seek between artists placing false cues and machines trying to classify artistic dialogue.

A basis for knowledge based media rich content - Enhanced Multimedia Meta Objects

One of the problems with combining knowledge mark-up with any standardised data models is that the data models of existing standards have insufficient modelling support for knowledge structures. With enhanced multimedia meta objects (EMMOs) CULTOS proposes an object structure in which certain object components can be described with standards such as MPEG-7 and tied together with (e.g. knowledge based) information that is outside the standard. We believe that the more technology moves towards a knowledge based economy, the more virtual objects will face the problem that they can no longer be viewed as either a catalogue record (Dublin Core), or a multimedia asset (MPEG-7) or a digital item that needs to be cleared by a rights management system (indecs), but such objects will appear as one or the other to different systems, depending on what is supposed to happen at each point in a value chain. We argue that EMMOs are a first step towards what could be called a "Smart Content Infrastructure". This paper only introduces the concept as it is central to the way in which

CULTOS operates. The idea is explained in more detail in (Goyal et al, 2002) and (Schellner et al, 2003).

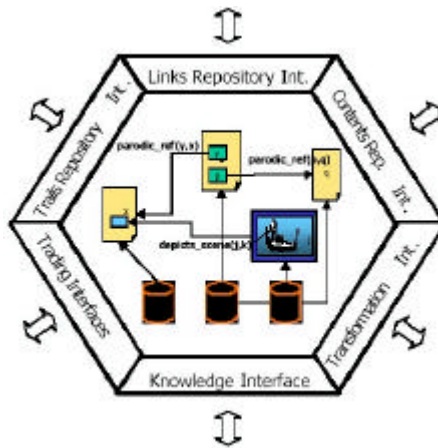


Figure 1 - An EMMO with its interfaces to standardised platforms

Lesson One - Usability

Our user group are experts in the humanities. They are frequent computer users but have a limited understanding of the underlying concepts such as databases, object hierarchies, or client server systems. However, they have high expectations concerning software usability and reliability. It is extremely difficult for a research project to satisfy the sum of expectations ranging from multimedia playing performance "as seen on a Realplayer", to "Undo" functionality as known from Word in conjunction with an object-oriented database model, to automated Webpage generation "as seen on flashy WWW sites". For the releases of the authoring tool, Salzburg Research essentially ran a hot-line as part of the validation phase. We believe this is necessary when one is not just building "tools for other computer scientists", but we had not accounted for it sufficiently in the budget and the following question needs to be answered: shall we leave this to product development, at the price of never even getting close to a product in these projects (and therefore never understanding the issues) or should we consider funding usability studies at even more than the 50% rate, because this is where the difference between uptake and losing out arises.

Lesson Two - Ontologies are the product of community consensus

At the beginning of the project, the idea of creating an ontology was an alien concept to several of the end user groups. As the project got under way, a lead group emerged who did virtually all of the ontology building with the other groups critiquing and proposing new concepts. One of the major changes from the original plan was the fact that instead of taking up the planned route of ontology specialisations for different purposes, the content group decided to actually agree on a single model, but do the modelling in such a way that different interpretations are made explicit through the ontology, rather than through building a variant of the ontology. On the other hand, it did become evident early on that local specialisations would be necessary so that thread authors could extend the ontology locally, within the given framework. The current ontology has approximately 300 intertextual relationships defined. Only the main group who developed the ontology and therefore also understood the intricacies of the knowledge modelling tool used, has a thorough capability of using all the constructs available in the ontology. In order for the rest of the ontology group to keep up to date, a workshop on how to use the ontology was necessary and it was also necessary for the lead group to supply a comprehensive glossary of the relationships and how to use them. All in all, three views on the ontology were created: firstly, a paper model of approximately one hundred concepts and relationships (Month 6). Secondly, a first implementation using the knowledge building tool supplied to CULTOS by one of the technical partners. Thirdly, the final glossary of 300 intertextual relationships was written that also plays the role of a proposal for a first ontology in the domain of intertextual studies (Benari & Ben Porat 2003).

Lesson Three - Build quality and formal rigour of ontologies

We found that apart from the domain understanding of the experts, three external factors influence the quality and usefulness of a domain ontology. Firstly, the knowledge building tool and the underlying knowledge representation mechanism influenced the building of the ontology. We found that the tool provided by one of the partners was user friendly and did certainly contribute to the successful

development of the ontology. The tool did have two weaknesses which were the lack of a standard knowledge representation model (although it can be classified as an object-oriented database language) and the lack of export options into any of the semantic web standards (XML Schema, RDFS, OWL). This was due to the system previously being used in ad-hoc knowledge modelling (informal semantic nets).

The second quality factor is the methodology used for creating the ontology. Unfortunately, we only became aware of the Ontoclean work (Guarino and Welty, 2002) after a significant part of the initial modelling had already been done. A re-assessment of our ontology in terms of the Ontoclean meta-modelling constructs would probably be beneficial. Similarly, the use of a common upper ontology would be desirable.

The third quality factor is the level of formality that is associated with the ontology in question. In the case of intertextual studies, the main aim has been to enable common practice through the use of a controlled vocabulary. Our ontology, albeit formally implemented, is not used by any inference engine to derive further knowledge from knowledge being put in explicitly. For serious semantic web applications, people will have to be able to pick the necessary level of formality before they start the project.

Lesson Four - Semantic web must not create WOKs - Write-Only Knowledge Systems

Having created multimedia threads with knowledge annotations and having stored them in a knowledge-enabled repository, we must ask ourselves what we can do with the knowledge thus gathered?

We postulated and are now in the process of developing, EMMO transformers which take the information structure and media objects and render them (possibly after some additional, presentation related authoring of the EMMO) on a web browser, or print the information in some linear format.

In a future project, we want to look at the question of how knowledge is communicated between agents and/or humans from a more fundamental angle: Can we - given a multimedia knowledge structure - get agents to interpret that structure, decide on a plan for a narrative, and have the narrative constructed? There may be additional authoring steps needed to embed enough hooks in the EMMO so as to allow for this.

Lesson Five - Organisation and Infrastructure of a Knowledge Based Economy

Let us finally look ahead into the future of CULTOS and its brothers and sisters. We are working towards a proposed politically announced) "knowledge based economy". Will the technologies be the bottleneck? From our limited experience in CULTOS we suggest that it will be the regulatory basis of the knowledge economy (i.e. a knowledge-ignorant regulatory basis!) that will pose problems. We have no satisfying way of providing access rights to people who would like to view EMMOs. And if we did have simple ways of access then we would be hampered by the lack of infrastructure. Rights clearance and viewer systems would have to work seamlessly in order to ensure customer satisfaction. Lack of interoperation and transparency for the user will delay the uptake of technologies.

Acknowledgements

In this brief paper we have presented the lessons learned in a European project, which aims at supporting literature professionals with knowledge-based multimedia authoring and presentation tools. The CULTOS project is part-funded by the European Commission, IST Programme, Knowledge Management and Content Creation.

References

Motti Benari, Ziva Ben Porat et al, Proposal for a Standard Ontology of Intertextual Studies, CULTOS Public Deliverable D3.3.2, June 2003.

Susan Blackmore, The power of MEMES. Scientific American, 280: 52–61, October 2000.

Guarino, N. and Welty, C.: Evaluating Ontological Decisions with OntoClean. Communications of the ACM, 45(2) (2002) 61-65.

Sunil Goyal and Wernher Behrendt and Siegfried Reich, EMMOs - Enhanced Multimedia Meta Objects for Re-purposing of Knowledge Assets, in MIS '02 - Meta Informatics Symposium 2002, Esbjerg, August 2002, pp155-160, LNCS 2641.

Karin Schellner, Utz Westermann, Sonja Zillner, Wolfgang Klas, "CULTOS: Towards a World-Wide Digital Collection of Exchangeable Units of Multimedia Content for Intertextual Studies", 2003 International Conference of Distributed Multimedia Systems, September 24-26, 2003.