

A Mini-Experiment in Semantic Annotation

A. Th. Schreiber, I.I. Blok, D. Carlier, W. P. C. van Gent, J. Hokstam, and U. Roos

University of Amsterdam, Social Science Informatics
Roetersstraat 15, NL 1018 WB Amsterdam, The Netherlands
E-mail: schreiber@swi.psy.uva.nl
WWW home page: <http://www.swi.psy.uva.nl>

Abstract. This paper describes a mini-experiment in using a tool for semantic annotation to index photographs of Windsor chairs, a type of antique furniture. The annotation tool makes use of an ontology based on art standards. We report on the experiences of subjects using the tool. The results suggest that a certain level of domain expertise is needed for semantic annotations, but also that an annotation tool has a clear added value for indexers.

1 Introduction

The semantic web requires semantic annotations of resources. The majority of the semantic-web research is at the moment targeted at languages (e.g., RDF(S), DAML+OIL) and architectures for the specification of semantic annotations. In our view however, there is also a real need to explore the issues related to the use of these techniques in applications.

We report on a small study in which subjects were asked to annotate objects with the help of an ontology. As subject domain we chose photographs of antique furniture, in particular Windsor chairs. In general, annotation and search for large sets of art-object images is becoming an important issue, as musea start making their collections available through the web. Within the art domain large knowledge corpora have been developed, such as the Art and Architecture Thesaurus [1], which makes semantic annotation in principle feasible. Another characteristic of the art-image domain is the manual indexing is standard practice. Musea routinely employ people to do this.

Not much is known yet about the semantic annotation process. The goal of this study is to get data about the following questions:

- What is the required level of domain expertise to be able to make useful semantic annotations?
- Can an ontology be presented in a way that is understandable to the indexers?
- What kind of tool support is useful for semantic annotation?

In Sec. 2 we briefly describe the semantic-annotation tool we used for this study and the ontologies which serve as the basis for the annotations. Sec. 3 describes the setup of the annotation study. In Sec. 4 we describe the results of the study. Finally, we discuss some issues arising from the study and speculate on the potential value of the semantic annotation for improving the search process.

2 Ontologies and Annotation Tool

For this study we used an ontology-based annotation tool that was previously developed to annotate photographs of apes [2]. In a nutshell, the tool reads in a set of ontologies that provide:

- a object-description template, i.e. a structured set of object properties, and
- subsumption hierarchies of concepts which act as potential property values for the object-description template.

The annotation tool assumes the ontologies are represented in RDF Schema. From the ontologies the tool automatically generates an annotation interface. The internal structure of the object-description template (e.g., grouping of properties) is represented in the tool interface through a tab structure (for an example, see the tabs in Fig. 1).

The ontologies for antique furniture we used in this study are described in more detail in a previous publication [3]. The basis of the object-description template is formed by the Visual Resource Association (VRA) Core Categories. VRA version 3.0 [4] defines 17 “data elements” for describing visual resources (with a bias towards the art domain). The VRA data elements are defined as a specialization of Dublin Core. In addition to the VRA data elements the description template contains eight data elements resulting from the European GRASP project. This project developed an ontology for describing and retrieving stolen art objects.

The data elements (i.e. properties of an antique-furniture object) were grouped into four sets:

1. *General features* such as the main type and the title of the object
2. *Production-related descriptors* such as creator (“maker”), style/period, technique, culture
3. *Physical descriptors* such as measurements, color
4. *Administrative descriptors* such as collection ID, rights, current location

In addition, we used the Art and Architecture Thesaurus (AAT) [1] as an ontology of concepts for describing art objects. Where possible, the value set of object-description properties is linked to relevant subhierarchies of AAT. For example, the property style/period is linked to three parts of the AAT hierarchy in which the appropriate style and period concepts can be found. For a detailed discussion of the underlying ontologies, of the RDF Schema representation and of the ontology mapping, the reader is referred to the aforementioned publication [3]. For this study we made one amendment: we restricted the main type of the art object to “Windsor chair”. AAT contains a hierarchy of Windsor chair types (in total 18 classes).

A snapshot of the annotation tool is shown in Fig. 1. The tab with “production-related descriptors” is shown, which contains six properties. The properties labeled with a “magnifier” icon (style/period, material and culture) are linked to hierarchies of AAT concepts.

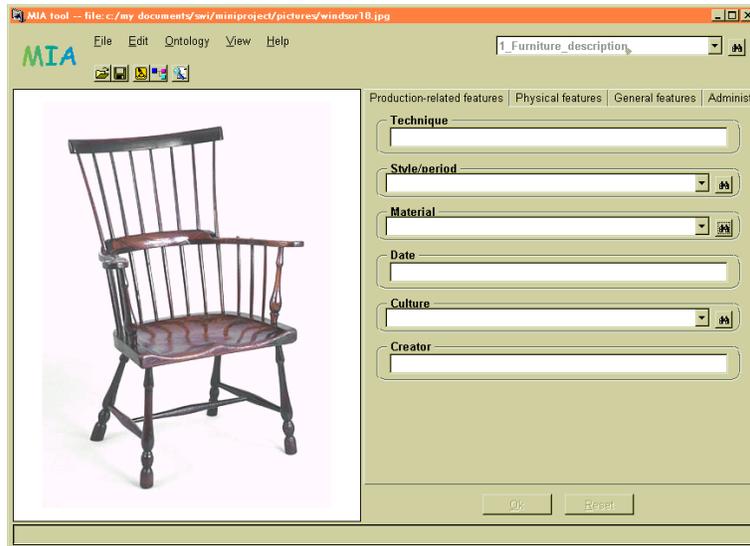


Fig. 1. Snapshot of the annotation tool for Windsor chairs. At the right, the tab for “production-related descriptors” is displayed

3 Annotation Study Setup

Subjects For this study we asked two “experts” and two “non-experts” to annotate the photographs. The experts were art historians working for a major auction house in Amsterdam. It should be noted that they were not experts on Windsor chairs, although they had some basic knowledge about it. The non-experts were information-science students.

Data collection Data collection consisted of four steps:

1. The subjects were asked to fill in a questionnaire containing in total eight questions about motivation, knowledge about the domain and computer/internet skills.
2. The non-expert subjects were given some background material on Windsor chairs and were asked to study this for 15 minutes. During the actual annotations they were allowed to consult this material. All subjects received a short oral instruction on the use of the annotation tool plus a one-page written instruction.
3. All subjects were presented with eight photographs of Windsor chairs and were asked to create an annotation for the image with the help of the tool. The first three photo’s were used as a training set. The following process data were collected while the subjects annotated the second set of five photographs:
 - (a) Time required per annotation object.
 - (b) Remarks made by the subjects during annotation.
 - (c) Resulting annotations (stored by the tool).
4. After annotating the photographs the subjects were asked to fill in a questionnaire about their experiences.

4 Results

Motivation, background knowledge and computer skills The two experts were interested in Windsor chairs. They had some basic knowledge of Windsor chairs, but were not experts in this particular subfield. The non-experts had had no prior interest in antique furniture. Both had heard about “Rococo chairs”, one about Windsor chairs, but they did not know exactly the meaning of these terms. All subjects worked with computers and internet browsers on a daily basis.

Annotation process Table. 1 lists the respective times required by the subjects to annotate the photographs. The annotation time decreases for successive images, which suggests a learning effect. The experts use considerably more time for annotations than the non-experts (34 and 30 minutes versus 15 and 21 minutes). During the annotation

Subject/ time	P1	P2	P3	P4	P5	Average time	Total time
Expert 1	10	6	5	7	6	6.8	34
Expert 2	8	9	5	4	4	6.0	30
Non-expert 1	4	3	3	3	2	3.0	15
Non-expert 2	6	5	4	3	3	4.2	21

Table 1. Time in minutes required by the subjects for annotating five photographs of Windsor chairs (P1-P5), plus the average and total time per subject

process, the experts said they found “the hierarchy” to be large and easy to understand. They made some remarks about the naming or ordering of properties, e.g. “First title and then type is more logical for us” and “Source is a bad name for origin”. Some features of the photo were difficult to establish, given the available information. One expert said: “We could use such a tool. There is nothing like this yet.” The non-experts mainly made remarks about the functionality of the tool (“how can/should I ...”) and about the meaning of terms.

A review of the resulting annotations made clear that (as expected) the annotations by the experts are more detailed. They fill in more slot values of the object-description template and the slot values themselves are more specific (i.e., at a lower level of the AAT hierarchy).

In the post questionnaire, all subjects agreed with the statement “this is a comfortable annotation tool”, and with the statement “the structure of the hierarchies is easy to understand”. Two experts and one non-expert agreed with the statement “the photo properties form a logical set”; one non-expert disagreed. The two experts disagreed with the statement “it is difficult to make annotation errors in this way”; one non-expert agreed and the other was unsure. All subjects disagreed with the statement “most people would annotate the object in the same way as I did” and with the statement “One does not need much background knowledge to annotate such photographs”.

5 Discussion

This small study does not warrant any strong conclusions. Some remarks can be made, however, on the questions raised in the introduction.

We were pleasantly surprised with the detailed comments made by the experts. They had no detailed knowledge about particular subject of Windsor chairs and seemed to use the ontology as a kind of on-line knowledge base. They clearly liked working with the tool and spent much time on the annotations. Such an annotation device surely is of potential value to them. The non-experts probably did not have enough domain knowledge to do the annotation job, but were not dissatisfied with the tool.

All subjects said that they expected other people might create different annotations. The damaging effect of this is probably less than in a pure keyword approach, due to the fact that some non-defined properties can be inferred. Also, if an indexer, e.g. due to lack of domain knowledge, annotates an object one level higher in the hierarchy, this only reduces the value of the annotation in a limited manner.

The subjects did not have many problems in understanding the hierarchy. The ontology structure makes clear which AAT concepts can be used for particular object features, which is likely to have helped here. The user interface of the tool can be improved, though. In the current tool, the interface is generated automatically from the ontology. This leads to non-logical ordering of tabs and of properties on a tab (see the remarks by the experts). It would have been better to introduce an explicit ontology-interface mapping into the tool, similar to the mechanisms in Protégé-2000 [5].

To get some (speculative) ideas about the added value of the semantic annotations for search purposes, we looked at a website for antique furniture, AntiqueArts.com, with in total 50K images. For general categories such as “chair” and “Windsor chair” retrieval is good (precision and recall > 0.8). However, for queries like “wooden chair” and “British chair”, recall of the search engine drops to almost zero. Semantic annotation would have enabled inferences to improve this drastically.

Acknowledgments This first author was supported by the ICES-KIS project “Multimedia Information Analysis” funded by the Dutch government.

References

1. Peterson, T.: Introduction to the Art and Architecture Thesaurus. Oxford University Press (1994) See also: <http://shiva.pub.getty.edu>.
2. Schreiber, A.T., Dubbeldam, B., Wielemaker, J., Wielinga, B.J.: Ontology-based photo annotation. *IEEE Intelligent Systems* **16** (2001) 66–74
3. Wielinga, B.J., Schreiber, A.T., Wielemaker, J., Sandberg, J.A.C.: From thesaurus to ontology. In Gil, Y., Musen, M., Shavlik, J., eds.: Proceedings 1st International Conference on Knowledge Capture, Victoria, Canada, New York, ACM Press (2001) 194–201
4. Committee, V.R.A.S.: VRA core categories, version 3.0. Technical report, Visual Resources Association (2000) URL: <http://www.gsd.harvard.edu/staffaw3/vra/vracore3.htm>.
5. Grosso, W.E., Eriksson, H., Fergerson, R.W., Gennari, J.H., Tu, S.W., Musen, M.A.: Knowledge modeling at the millennium: The design and evolution of Protégé-2000. In: 12th Banff Workshop on Knowledge Acquisition, Modeling, and Management. Banff, Alberta. (1999)