Evaluating DILIGENT Ontology Engineering in a Legal Case Study

Christoph Tempich[1], H. Sofia Pinto[2], York Sure[1], Denny Vrandečić[1], Núria Casellas[3], and Pompeu Casanovas[3]

[1]Institute AIFB, University of Karlsruhe, Karlsruhe, Germany
{tempich,sure,vrandecic}@aifb.uni-karlsruhe.de
[2]Dep. de Engenharia Informática, Instituto Superior Técnico, Lisboa, Portugal
sofia.pinto@dei.ist.utl.pt
[3]Institute for Law and Technology (IDT). Universitat Autonóma de Barcelona, Spain
{nuria.casellas,pompeu.casanovas}@uab.es

Abstract. In the IST project SEKT – Semantically Enabled Knowledge Technologies – one of the case studies [BCC+05] aims at providing an intelligent Frequently Asked Questions system, *Iuriservice prototype II*, that will offer help to newly appointed judges in Spain. This iFAQ uses ontology based Semantic Web technologies to retrieve FAQ to reply queries posed in natural language. In order to construct the underlying ontology, we used the ontology engineering methodology DILIGENT.
DILIGENT [PTSS04] methodology is intended to support domain experts in a distributed setting to engineer and evolve ontologies with the help of a fine-grained methodological approach based on Rhetorical Structure Theory [MT87]. Although, knowledge acquisition and other content issues are not discussed in this paper, harmonization efforts for knowledge modelling are required by discussing the ontology. During the discussion, participants exchange arguments which may support or object to certain ontology engineering decisions. As any legal professional knows, tracking the argumentation of decisions is essential in allowing later readers to understand them.
We will present an evaluation framework to evaluate the methodology within this particular context. We will take a detailed look at the development of the ontology and especially the novel experience of using a wiki to support the collaborative ontology engineering and to capture the ongoing discussion of the ontology. First results of this evaluation will be presented.

Keywords: Ontology Engineering, methodology, Rhetorical Structure Theory (RST)

1 Introduction

Ontologies in the legal domain have a long history, but they usually focus on formalising the normative knowledge codified in laws and judicial decisions. The legal case study of the EU IST project SEKT – Semantically Enabled Knowledge

Technologies – differs from that radically, as it tries not to model the normative knowledge but rather the much more diverse knowledge of legal professionals, especially judges, in their daily work. We will begin with a brief description of the SEKT legal case study in section 2.

Next we describe the initial processes used in developing the ontology, and the problems encountered, in section 3. The ontology engineering team focused on the competency questions used for building the ontology. We describe how they were applied, followed by the introduction of the DILIGENT argumentation framework and the tracking of arguments in section 4. These changes introduce parts of the DILIGENT ontology engineering processes, which is described in much more detail in [PTSS04].

Although DILIGENT was already evaluated in different settings (refer, for example, to [PSST04] for a case study), the legal case study offers a new and fertile setting for further evaluation, due to its setting in the legal domain, which is novel for the developers of DILIGENT. In section 5 we will evaluate the changes introduced to the ontology engineering processes, and highlight the positive implications and the deficiencies of the processes applied. Based on this, we give a short outlook to the further development of the methodology and the tools supporting it in the last section.

2 The Iuriservice II scenario

The goal of the legal case study is to provide support to professional judges. In the Spanish system one particular problem young judges face is when they are on duty and confronted with situations in which they are not sure what to do. In such cases, they usually phone their former training tutor (experienced judges) for resolving the issue. But this is a slow and insecure procedure. In this case study, it is planned to develop an intelligent system to speed up the process and to relieve experienced judges from this effort by providing support to young judges. Only in the case that the requested knowledge is not in the system and cannot be reformulated from already stored knowledge, an experienced judge will be contacted. The result of this "expensive" consultation will be fed back into the system automatically.

In order to build a scalable and useful system, several requirements have been identified [RCCP04]. In particular, the decision as to whether a request for knowledge is covered by the knowledge stored in the system should be based on semantics of the legal domain rather than on simple word matching. An ontology can be used to perform this semantic matching. Case-based reasoning techniques will be used when considered applicable. The main difference with traditional CBR approaches is that we will use a semantic-based similarity measures based on ontologies. Moreover, cases will be - where possible - automatically extracted from information generated by judges as they perform their daily work.

The ontology used for this semantic matching, the OPJK (Ontology for Professional Judicial Knowledge) [BCC+05], is being developed in the SEKT project. For further information on the OPJK refer to [CCP+05], here we will

present the problems that have arisen during the development of the ontology and how different approaches towards solving these problems have worked out.

3 The initial processes

In order to build the ontology, the UAB legal experts team first collected and consolidated a list of over 800 competency questions that were expected to be answered by the *Iuriservice II prototype*. These questions were collected in a strenuous and lengthy process. In order to process the vast number of competency questions, they were first clustered in different groups, so that the ontology engineering process could be steered by the themes identified in the questions.

Using the competency questions, the relevant terms and relations in the questions were identified. The team then discussed the need to represent them within the ontology and their place within the taxonomy. And accordingly, we followed the middle-out strategy [GPFLC03]. With this strategy, the core of basic terms are identified first and then they are specified and generalized if necessary. Finally, the relevant relations between those terms also have to be identified (mainly *is a* and *instance of*). Furthermore, some other relations were also identified: someone creates a *document*, thus there is a relation *has author*.

One of the main hindrances when building the ontology was the wish to represent the answers within the ontology. Only by a closer analysis of the planned usage of the ontology within the *Iuriservice II* system it was realised that the ontology will not be used to answer the question but rather to allow for a semantic matching of the naturally poised questions. Thus, instead of modelling the answers within the ontology, modelling the questions was enough. This reduced the complexity of the ontology drastically and especially allowed for a well-defined limit inside the discussions.

However, difficulties in reaching consensual decisions and the lack of traceable lines of argumentation was slowing down the construction of the ontology. For that reason, the introduction of the Distributed, Loosely-controlled and evolving Engineering of oNTologies (DILIGENT), provided by the AIFB research team, offered a reliable basis for a controlled discussion of the arguments in favour and against modelling decision. The introduction of DILIGENT proved the need to rely on guidelines for the decision-making process within ontology design.

4 DILIGENT Argumentation Framework and tracking

There is evidence that distributed ontology development can be rather time consuming, complex and difficult, in particular getting agreement among domain experts. Therefore, one needs an appropriate framework to assure it in a speedier and easier way. In order to provide better support, one needs to identify which kind of arguments are more relevant and effective to reach consensus. The Rhetorical Structure Theory (RST, see [MT87]) can be used to classify the kinds of arguments most often used and identify the most effective ones. The ontology engineering methodology DILIGENT [PTSS04] is intended to support domain

experts in a distributed setting to engineer and evolve ontologies with the help of a fine-grained methodological approach based on RST.

With the help of RST and case studies we identified the the most effective arguments and created an argumentation ontology described in [TPSS05]. The discussions that came up when building the ontology were focussed by allowing to concentrate only on the given set of arguments, and thus shortened much.

The arguments exchanged during the discussion can also serve as a rationale about the ontology engineering decisions. Tracking these arguments solves many of the problems described in the previous section: the tracked arguments allow anyone to return to the discussion and understand the reasons for a certain decision, both for the new user who entered the engineering team later as well as the users who were there at the decision itself, but can not remember the rationale behind it. As any legal professional will agree, tracking the argumentation behind decisions is essential in order to sustain a consistent system.

We offered an easily accessible web based interface in order to allow the discussion in a traceable way. A standard wiki was used which supports seamless discussion and offers ease of use. But the users quickly extended it in combination with the KAON OIModeller [GSV04]: they modelled the ontologies agreed on, made a snapshot of part of ontology and imported it to the wiki, in order to visualize the ontology and ease the discussion and understanding of it. In figure 1 we see a screenshot of the wiki, running on the SEKT portal, showing the concept HECHO with its description, argumentation and attributes, as well as a graph made with the KAON OIModeller showing the concept and its attributes.

5 Evaluation of the processes

In the case study two ontology engineering processes were used. The initial one was used for about five months and included a lot of discussions, but finally all results were abandoned. The second process, following DILIGENT, was started afterwards. In only a few weeks time the UAB legal experts team managed to pass through 200 of the competency questions and to agree on an ontology with 50 concepts, 100 relations and over 300 instances (following the explicit goal to keep the number of concepts small).

Geographically distributed members of the team were finally able to follow the discussions, due to the capture of the argumentation on the web-accessible wiki. Also new members could refer to wiki in order to understand the rationale of certain modelling decisions. Before the introduction of the wiki this was not possible: the only possibility to get these informations was by asking the members of the team that were involved in the discussion, and hoping that they remember.

Using thematically clustered competency questions instead of proccessing them in an unordered way brought a noticeable speed up: there was no need to keep changing the context from question to question. The discussion was more efficient by keeping the cognitive distance of one question to the next small.

As there are no protocols of the discussions during the initial ontology engineering process, we cannot compare the effect of the introduction of the RST

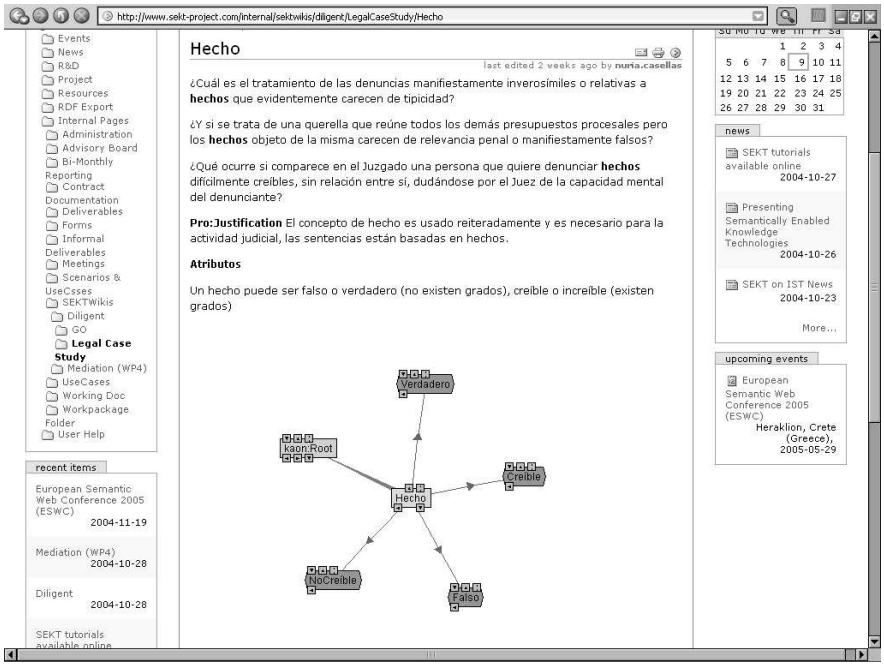

Fig. 1. Screenshot of the wiki used for discussing the legal ontology

based DILIGENT Argumentation Framework. Although the members of the UAB legal experts team claim a higher efficiency in their discussions, this can also stem from other sources (like having the former discussion available, or building on an already agreed-on foundation of a few concepts). We have to rely on former case studies (like [PSST04]) for this question.

6 Conclusion

As we have seen in the previous section, we gained quite some efficiency by applying DILIGENT. DILIGENT improves the ontology engineering process, by following a principle that is accepted for a long time in the legal domain: traceability of argumentations. But we also found some deficiencies with applying our methodology that we have to investigate further.

For the tracking of the ontology we used an off the shelf standard wiki. But the UAB legal experts team demands graphical, intuitive and easy accessible user interfaces. They want tools that help them with the specific tasks they encounter when dealing with ontology engineering. They have gone great lengths to combine the graphical ability of the KAON OI modeller with the flexibility and easy accessibility of the web based wiki software.

So the given tool support as it is now is not optimal. A standard wiki can only be one building block of a DILIGENT inspired ontology engineering tool.

The users in the given case study were highly motivated and expanded the usage of the tool creatively in order to achieve what they wanted, but we cannot expect this from all users.

Nevertheless, the evaluation of DILIGENT was very positive. We achieved big gains in terms of speed of development and documentation of the results. We will continue to investigate in the above mentioned issues in order to enhance DILIGENTand the tools associated with the methodology.

Acknowledgements. Research reported in this paper has been financed by EU in the the IST project SEKT (IST-2003-506826).

References

[BCC+05] V.R. Benjamins, P. Casanovas, J. Contreras, J.M. Lopez Cobo, and L. Lemus. Iuriservice: An intelligent frequently asked questions system to assist newly appointed judges. In V.R. Benjamins, P. Casanovas, A. Gangemi, and B. Selic, editors, *Law and the Semantic Web*, LNCS, Berlin Heidelberg, 2005. Springer.

[CCP+05] P. Casanovas, N. Casellas, M. Poblet, J. Vallbé, Y. Sure, and D. Vrandečić. Iuriservice ii ontology development. In *Workshop on Artificial Intelligence and Law at the XXIII. World Conference of Philosophy of Law and Social Philosophy*, 2005.

[GPFLC03] A. Gómez-Pérez, M. Fernández-López, and O. Corcho. *Ontological Engineering*. Advanced Information and Knowlege Processing. Springer, 2003.

[GSV04] T. Gabel, Y. Sure, and J. Voelker. KAON – ontology management infrastructure. SEKT informal deliverable 3.1.1.a, Institute AIFB, University of Karlsruhe, 2004.

[MT87] William C. Mann and Sandra A. Thompson. Rhetorical structure theory: A theory of text organization. In Livia Polanyi, editor, *The Structure of Discourse*. Ablex Publishing Corp., Norwood, N.J., 1987.

[PSST04] H. S. Pinto, S. Staab, Y. Sure, and C. Tempich. OntoEdit Empowering SWAP: a Case Study in Supporting DIstributed, Loosely-Controlled and evolvInG Engineering of oNTologies (DILIGENT). In C. Bussler et al., editors, *Proceedings of the 1st ESWS 2004*, 2004.

[PTSS04] H. Sofia Pinto, Christoph Tempich, Steffen Staab, and York Sure. Diligent: Towards a fine-grained methodology for distributed, loosely-controlled and evolving engingeering of ontologies. In Ramon López de Mántaras and Lorenza Saitta, editors, *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI 2004), August 22nd - 27th*, pages 393–397, Valencia, Spain, AUG 2004. IOS Press.

[RCCP04] L. Rodrigo, M. Blázquez Cvico, P. Casanovas, and M. Poblet. Legal case study before analysis. SEKT deliverable 10.1.1, Intelligent Software Components S.A. and Universitat Autonoma de Barcelona, 2004.

[TPSS05] Christoph Tempich, H. Sofia Pinto, York Sure, and Steffen Staab. An argumentation ontology for distributed, loosely-controlled and evolving engineering processes of ontologies (diligent). In C. Bussler, J. Davies, D. Fensel, and R. Studer, editors, *Second European Semantic Web Conference, ESWC 2005*, volume YYYY of *LNCS*, pages xx–yy, Heraklion, Crete, Greece, MAY 2005. Springer. accepted.