

ONTORULE Annual Report

ONTORULE



www.ontorule-project.eu

The objective of ONTORULE is to enable users, from business executives over business analysts to IT developers, to interact in their own way with the part of a business application that is relevant to them.

We believe that one essential step towards achieving that objective is the ability to separate cleanly the domain ontology from the actual business rules, on the one hand; and the representation of the knowledge from its operationalization in IT applications, on the other hand.

Leading vendors of knowledge-based and business rules management systems and top research institutions join their efforts, in ONTORULE, to develop the technology that will empower business professionals in the knowledge economy of the future.

Two large industrial companies are the test beds that will ensure the success and business impact of the technology.

Summary of Activities

The first year of the project has been devoted to building firm foundations for ONTORULE. Beside the time and effort given to setting up an effective and efficient communication and collaboration organization, training each other in the project's many background technologies, and researching the state of the art (especially with respect to the combination of ontology-based and rule-based inference), most of the effort was spent on three major objectives:

- The analysis of the use cases' requirements. The effort to analyse the requirements of our two industrial use cases in the perspective of the ONTORULE vision initiated our work on articulating the ONTORULE architecture;
- The implementation of bottom-line technology for ontology acquisition from business experts and textual policies, the operationalization of the acquired knowledge, and the combination of rules and ontologies for inferencing purposes;
- The design and specification of foundational standards. The ONTORULE team has been instrumental in the emergence of new standards at OMG and W3C.

Next year, building on these foundations, we expect to develop the first effective version of the ONTORULE platform, and to implement the initial prototypes of the industrial use case applications.

Use case requirements and architecture

The BPMN diagrams, below, show a (simplified) view, in an architecture-like form, of the two industrial use cases: ArcelorMittal's steel use case is about assigning a destination to steel coils, based on their characteristics and the order specification; and AUDI's CAx use case is about the semantic integration of CAD (computer aided design), CAE (computer aided

engineering) and CAT (computer aided testing) tools. The pilot application will focus on implementing test requests.




Figure 1. Simplified BPMN diagram of the ArcelorMittal use case




Figure 2. Simplified BPMN diagram of the AUDI use case

Bottom-line technology

From a technology point of view, the ONTORULE vision is that, based on an effective methodology to involve the business user, and starting from the original sources of a policy (including documents in their native form – that is, usually, mostly natural language, sometimes mixed with decision tables):

1. a domain ontology and a domain terminology will be built with active participation of the business experts; and, based on the ontology and the terminology, and using the

OMG SBVR (semantics for business vocabularies and business rules) meta-model as a guidance, the business rules that are specified in the source will be extracted, with the active participation of the business experts as well, and represented in a structured and formal language;

2. the acquired knowledge will be translated in the operational languages most appropriate for storage and interchange; for management and maintenance; and for use in decision support applications, where it is unlikely that one single language will be the best choice, or even usable, for the different categories of knowledge and for the different categories of operations;
3. The knowledge will be used to produce inferences, first and foremost to support decision making by the end-user, but also to support the owner of the knowledge with respect to knowledge management, maintenance and evolution.

During this first year of ONTORULE, we have developed prototypes covering all three aspects. These prototypes are basic, and they are intended to serve as the basis for further research, exploration and experimentation. The ONTORULE team has, thus, developed initial prototypes for (the list is not exhaustive):

- the interactive extraction of terminology and ontology from natural language policy texts;
- the translation of SBVR models of business rules and business vocabularies into operational languages such as W3C OWL and IBM Websphere ILOG JRules's rule language;
- a production rule inference engine that delegates most of the filtering to a description logic inference engine, thus allowing production rules to be grounded directly in W3C OWL ontologies;
- an inference engine for a logic language that combines F-Logic and description logics.

Foundational standards

Standards are at the core of the ONTORULE vision and they are bound to be everywhere in the ONTORULE architecture: the main deliverable of the ONTORULE project, that is, the specification of the ONTORULE platform, will be specified entirely in terms of open standards.

It is, therefore, essential to the success of ONTORULE that existing standards that might be useful to the project's purpose be identified, as well as areas where ONTORULE will require interchange or interoperability amongst functions or modules, but no open standards exist.

To contribute to that purpose, ONTORULE publishes a periodic standards monitor.

Moreover, the approach chosen by ONTORULE is to engage actively in standardization efforts, even to take leadership positions, where open standards under development are most likely to be useful and where new standards will be required.

At this stage, members of the team play a prominent role in a number of on-going standardization efforts of high importance to the project. Team members act, most noticeably,

- as co-chair of the W3C rule interchange format (W3C RIF) working group, and as editors or co-editors of three specifications on the W3C RIF recommendation track, that are of particular significance to ONTORULE (*RIF production rule dialect*, *RIF combination with RDF and OWL* and *RIF combination with XML data*);
- as active members of the OMG SBVR revision task force.

ONTORULE members have also been actively monitoring or participating in the development of W3C OWL 2 (Web ontology language), SKOS (simple knowledge organization system) and SPARQL (SPARQL protocol and RDF query language), and of OMG BPMN 2.0.

User Involvement, Promotion and Awareness

ONTORULE partners have published and presented research papers related to ONTORULE in several occasions, including prestigious conferences such as the international joint

conference on artificial intelligence (IJCAI) or the European semantic Web conference (ESWC).

The project and its objective has also been featured prominently in presentations at the Business Rules Forum, the Business Rules Platform Netherlands, the OnTheMove conference and technical meetings of the object management group (OMG) and the world wide Web consortium (W3C).

ONTORULE has introduced itself to other projects at the 1st EU “matchmaking” event, which was organized by STI International, with the support of the European Commission, in collocation with ESWC.

The consortium aims at organising ONTORULE tutorials at major conferences, next year, such as the conference of the association for the advancement of artificial intelligence (AAAI), the European conference on Artificial Intelligence (ECAI), the world wide Web conference (WWW) and the international and European semantic Web conferences (ISWC and ESWC).

Future Work

The second year of the ONTORULE project will be a year of realization:

- a first version of AUDI's and ArcelorMittal's pilot applications will be developed, using the bottom line technology that we developed during the first year, for the purpose of providing feedback and of refining the applications requirements on ONTORULE;
- on the other hand, the bottom line will be raised to the point where the technology is actually usable, a first version of the specification of the ONTORULE architecture will be finalized, and we will experiment with initial prototypes of open standards-based integration.

On the standardization front, we aim at finalizing the current rounds of W3C and OMG specifications where ONTORULE members play a leading role (W3C RIF, OMG SBVR, OMG PRR, the production rule representation meta-model), and at initiating the next round where useful. We will, especially, watch the new initiative on decision modelling and notation, at the OMG; and work towards application-ready RIF dialects, building upon the foundations provided by the current set of basic specifications.

Watch for the demonstrators, on the project Web site: whereas the first year demonstrators are videos, because the technology is still fragmented and of complex use, you will be able to interact with next year's demonstrators and to use them!

Watch, also, for ONTORULE tutorials: we have submitted proposals to the AAAI, ECAI and WWW conferences in 2010 (AAAI-10, 24th AAAI conference on Artificial Intelligence, 11-15 July 2010, Atlanta, GA, USA; ECAI 2010, 19th European conference on Artificial Intelligence, 16-20 August 2010, Lisbon, Portugal; WWW10, 19th World Wide Web conference, April 26-30, Raleigh, NC, USA).

Further Information

Stay updated about ONTORULE progresses, on www.ontorule-project.eu.