

Context in the Large: Contextual Issues when Dealing with Thousands of Ontologies

Enrico Motta¹ and Mathieu d'Aquin¹

¹Knowledge Media Institute, The Open University, Milton Keynes, UK
e.motta@open.ac.uk, m.daquin@open.ac.uk

Abstract: Back in 2005 we launched a research programme called “Next Generation Semantic Web Applications”, which pioneered the vision of exploiting the Semantic Web as a source of background knowledge to enable the development of a new class of intelligent applications. This research programme was based on some rather radical tenets, envisaging applications relying not simply on a few hand-picked ontologies, but on the dynamic selection and use of knowledge acquired from thousands of online ontologies. Five years later, this programme has produced a variety of applications, tools and techniques, concretely realizing the paradigm in scenarios such as ontology evolution, relation discovery, word sense disambiguation, and semantic enrichment of folksonomies.

An important challenge when building applications that consume knowledge dynamically sourced from thousands of ontologies concerns dealing with the contextual nature of ontological resources. Specifically, new methods were needed to identify automatically whether the context of the original ontology still holds when the knowledge is reused in the application in hand.

In addition, we have also looked at this issue in an application-independent way and carried out analytical studies that consider the Semantic Web itself as an object of study and attempt to characterize and understand the relations between the different epistemologies which are at the basis of the published conceptualizations. In particular, we have developed formal notions of agreement and disagreement between ontologies and used these to cluster ontologies which seem to share similar world views. To our knowledge these studies represent the first ever empirical analysis of large scale distributed conceptualizations and provide useful insights into the concrete practices used by ontology engineers. Specifically, these studies make it possible for us to see which contextual viewpoints tend to occur more frequently and which communities share specific conceptualizations.

Finally, such results can be used further to develop highly structured ontology repositories, identifying and resolving contextual discrepancies between ontologies to facilitate and improve the efficiency of both automatic and manual access to resources on the Semantic Web.