

Exploring modelling choices: using non-functional characteristics to test an ontology add-on

Jennifer D. Warrender and Phillip Lord

School of Computing Science, Newcastle University, Newcastle-upon-Tyne, UK

Ontologies present an attractive technology for describing bio-medicine, because they can be shared, and have rich computational properties. However modelling choices affect the size, complexity and potentially the performance of an ontology. What we need is some systematic method of quantitatively testing an ontology to see what effect modelling choices have on the reasoning performance of the ontology.

There is relatively little work on addressing this form of non-functional characteristic testing. Once a large ontology has been built, it is often too difficult to change the form of modelling which is a requirement for understanding the performance impact; conversely, testing performance of small-scale or subsets of an ontology is often not indicative of performance scaling.

Taking a highly programmatic, and pattern-driven approach to ontology building means that it is possible to change and refine ontological models rapidly, even where these changes affect many entities within the ontology. We have pioneered this form of ontology building in the Karyotype ontology; an ontological implementation of the human cytogenetic nomenclature, using the Tawny-OWL software.

In this talk, we show a novel means to explore modelling choices and how these scale. We investigate the performance of several different axiomatisations of the same knowledge. We do this by extending the Karyotype Ontology. Specifically, we are investigating the best modelling representation of a deletion or inversion event where the event affects a sequence of bands between two breakpoints. We show that there are two viable ways of implementing the affects restriction; the enumeration of bands (using a generic closure axiom) and the implementation of a datatype restriction of ordinal bands, depending on the number of karyotypes we wish to model.

However, the non-functional scalability requirement is probably not the best basis to make for this design decision. Instead, with Tawny-OWL, we can make all implementations available and allow downstream users of the Karyotype Ontology to decide which implementation is best for them. Thus as future work, we could implement modelling choices through the use of contextual imports.