

The manual construction of ontologies is a hard and expensive task that usually requires knowledge of domain experts and skills of ontology engineers combined together. Therefore, methods that can automatically construct (parts of) ontologies are desirable. Such methods can use domain data as a knowledge source. As some domain knowledge can already be encoded as an (incomplete) ontology, those methods should also take it into account in order to produce informative additions.

DL-miner is a tool that automatically acquires, or mines, expressive schema-level knowledge from instance-level data in OWL. It implements a semantically faithful approach to the automated construction of ontologies and is aimed at acquiring maximally informative, non-contradictory axioms. More specifically, it respects already encoded domain knowledge and the standard OWL semantics with its open world assumption. The approach does not require human supervision and only asks a user for final judgements of produced axioms which are viewed as hypotheses. It mines general class inclusions as well as property inclusions. The tool orders hypotheses by several quality measures such that a user can sort and browse hypotheses according to their preferred quality criteria.

Our experiments show that mined hypotheses seem to be a good starting point for knowledge modellers. If some knowledge is already encoded, hypotheses appear to be informative missing bits. We also observe that DL-miner can reveal interesting relations in our data that can expose new knowledge about the domain and help to understand the data. In addition, hypotheses can contain “strange” axioms which may help us highlight, on the one hand, odd or erroneous modelling and, on the other hand, inaccurate or abnormal data. Thus, besides ontology engineering, DL-miner is potentially useful for semantic data analysis and ontology validation.