Abstract

A number of repositories exist containing a substantial number of ontologies, e.g. NCBO BioPortal and TONES. Following the principles of data mining, a natural question is: given such a repository, can we extract something that can be useful for a better understanding and design of ontologies. In this thesis, we introduce a strategy for mining a repository of ontologies in order to extract higher order knowledge patterns that can be used to refine, complete and classify ontologies; detect potential inconsistencies and suggest modeling alternatives; and hence investigate and improve the organization of the ontologies within the repository.