Quantification in Defeasible DLs matters again

Maximilian Pensel* and Anni-Yasmin Turhan

Institute for Theoretical Computer Science, Technische Universität Dresden,
first-name.last-name@tu-dresden.de

Defeasible DLs have nonmonotonic semantics and are a prominent approach for reasoning w.r.t. inconsistency or default assumptions. A *defeasible knowledge base* (DKB) contains *defeasible concept inclusions* (DCIs) in the so-called *DBox*. DCIs capture concept inclusions that *typically* hold and could be ignored during reasoning if they contradict more specific DCIs or the TBox. The more DCIs are consistent with an instance of a concept, the more *typical* it is considered. Rational reasoning as characterised by the well-known KLM postulates for propositional logic has previously been lifted to $\mathcal{ALC}$ by Casini et al. in 2010. The existing computation algorithms for entailments use a reduction to classical reasoning (by so-called materialisation) for rational entailments and for the stronger relevant entailments. Materialisation transforms each DCI $C \mathrel{\sqsubseteq\!\!\!\sim} D$ into a concept $\neg C \sqcup D$ which can be used for classical subsumption tests in conjunction with the (potential) subsumee to include the information from the DCIs. Obviously, such a conjunction does not propagate defeasible information to concepts nested in existential or value restrictions. Thus un-defeated DCIs are ignored, expected inferences are lost, and all role successors are of uniform (a-)typicality. —a well-known problem to the Defeasible DL community.

We introduced the first approach for rational reasoning in $\mathcal{EL}_\bot$ mending these shortcomings by extending (classical) canonical models to *typicality models*. Intuitively, the domain of a typicality model contains multiple copies of the domain of an $\mathcal{EL}_\bot$ canonical model. Each such copy satisfies a different (sub-)set of DCIs from the DKB. The copies allow to tailor the typicality of role successors required by, say $\exists r.C$, as follows: for $(d, e) \in r^\mathcal{I}$ and $e \in C^\mathcal{I}$ element d can be related to a more typical representative of C than e. Typicality models are computed by first performing classical reasoning w.r.t. a TBox and then successively "upgrading" typicality of role successors, i.e., picking more typical, yet consistent representatives of the (successor) concept. We showed that subsumption under rational closure is achieved, if the considered sets of DCIs are a particular *sequence of decreasing subsets* of the DBox. We have generalised our approach to the stronger relevant closure by using the *complete lattice of subsets* of the DBox for generating copies of the canonical domain. We show for rational and relevant closure that typicality models yield more entailments than the respective materialisation-based approach.

Our results for rational closure are published at the *14th International Conference on Logic Programming and Nonmonotonic Reasoning* (LPNMR 2017) and the results for relevant closure are published at the *4th International Workshop on Defeasible and Ampliative Reasoning* (DARe 2017).

* Supported by DFG in the Research Training Group QuantLA (GRK 1763).