Making Quantification Relevant Again —the Case of Defeasible $\mathcal{EL}_\bot$

Maximilian Pensel and Anni-Yasmin Turhan*

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

Abstract. Defeasible Description Logics (DDLs) extend Description Logics with defeasible concept inclusions. Reasoning in DDLs often employs rational or relevant closure according to the (propositional) KLM postulates. If in DDLs with quantification a defeasible subsumption relationship holds between concepts, this relationship might also hold if these concepts appear in existential restrictions. Such nested defeasible subsumption relationships were not detected by earlier reasoning algorithms—neither for rational nor relevant closure. Recently, we devised a new approach for $\mathcal{EL}_\bot$ that alleviates this problem for rational closure by the use of typicality models that extend classical canonical models by domain elements that individually satisfy any amount of consistent defeasible knowledge. In this paper we lift our approach to relevant closure and show that reasoning based on typicality models yields the missing entailments.

1 Introduction

Description Logics (DLs) are usually decidable fragments of First Order Logic. In DLs *concepts* describe groups of objects by means of other concepts (unary FOL predicates) and roles (binary relations). Such concepts can be related to other concepts as sub- and super-concepts in the TBox which is essentially a theory constraining the interpretation of the concepts. One of the main reasoning problems in DLs is to compute subsumption relationships between two given concepts, i.e., decide whether all instances of one concept must be necessarily instances of the other (w.r.t. the TBox).

While classical DLs allow only for monotonic reasoning, defeasible DLs admit a form of non-monotonic reasoning and have been intensively studied in the last years [5,6,7,3,4,8]. Most defeasible DLs allow to state relationships between concepts by defeasible concept inclusions (DCIs), which characterise typical instances of a concept and can be overwritten by more specific information that would otherwise cause an inconsistency. Often the semantics of defeasible DLs is

*This work is supported (in part) by the German Research Foundation (DFG) in the Research Training Group QuantLA (GRK 1763) and within the Collaborative Research Center SFB 912 – HAEC.

based on a translation of propositional preferential and (the stronger) rational reasoning for conditional knowledge bases introduced by Kraus, Lehmann and Magidor (KLM) in [10] to DLs. Recent studies on DDLs investigate different semantics, such as a typicality operator under preferential model semantics in [8], a syntactic materialisation-based approach [5,4], characterised with a different kind of preferential model semantics in [3], and extensions of rational reasoning to the inferentially stronger lexicographic and relevant closure in [4,7].

We consider here an extension of the lightweight DL $\mathcal{EL}$. In this DL complex concepts are built by conjunctions and existential restrictions, which are a form of quantification and clearly not expressible by propositional logic. The DL $\mathcal{EL}$ enjoys good computational properties: subsumption can be computed in polynomial time [2]. Despite its moderate expressivity, many applications rely on $\mathcal{EL}$, predominantly the bio-medical domain and the web ontology language with the OWL 2 EL profile. In contrast to $\mathcal{EL}$, its extension $\mathcal{EL}_\bot$ can express disjointness of concepts and thus inconsistencies. We consider in this paper non-monotonic subsumption under relevant closure in defeasible $\mathcal{EL}_\bot$.

In [5] Casini et al. showed that the complexity of non-monotonic subsumption coincides with the complexity of classical reasoning of the underlying DL and devise reasoning algorithms for defeasible subsumption under rational and relevant closure. Their algorithm uses a reduction to classical reasoning and thereby allows to employ highly optimised classical DL reasoners for the reasoning task. Their reduction uses materialisation, where the idea is to encode one consistent subset of the defeasible statements as a concept which is then used in the classical subsumption query as an additional constraint for the (potential) sub-concept in the query. Essentially, the algorithms for the two types of closure differ in the subsets of DCIs from the knowledge base they use for reasoning. While rational closure utilises only a single sequence of decreasing subsets of DCIs, the stronger relevant closure admits *any* such subset during reasoning. Thus relevant reasoning is done w.r.t. a lattice of DCI sets which include more combinations of DCIs potentially leading to more fine-grained entailments. However, both of the resulting algorithms in [4,5,7] are not complete in the sense that not all expected subsumption relationships are inferred. The reason is, that defeasible knowledge is not propagated to concepts nested in existential restrictions and thus even un-defeated knowledge is omitted during reasoning.

The goal of this paper is to devise a reduction algorithm for reasoning under relevant closure for $\mathcal{EL}_\bot$ that derives defeasible knowledge for concepts nested in existential restrictions. We have recently devised an approach that achieves this for reasoning under the weaker rational closure by the use of typicality models [12]. These models are an extension of the well-known canonical models for classical DLs of the $\mathcal{EL}$ family where the domain consists of elements representing the concepts occurring in the TBox. Now, typicality models have representatives for each pair of a concept occurring in the TBox and a set of defeasible statements. Thus, for the case of relevant closure such typicality models are built over a lattice-shaped domain according to the lattice of DCI subsets. For a simple form of these typicality models we show that it results in the same entailments as the materialisation-based approach [4] for relevant closure. We then extend the simple typicality models to remedy the mentioned shortcoming regarding existential restrictions. The main idea is, to admit in this kind of model differing "amounts" of consistent defeasible information for different occurrences of the same nested concept.

This paper is structured as follows: the next section introduces the basic notions of (D)DLs and $\mathcal{EL}_\bot$. Section 3 recalls the materialisation-based approach for rational and relevant closure and investigates their shortcomings. Section 4 introduces minimal typicality models over a lattice domain and shows that the same subsumption relationships under relevant closure can be inferred as by the materialisation-based approach from [4]. In Section 5 we extend these models to maximal typicality models over a lattice domain and show that these allow to obtain the formerly omitted entailments. The paper ends with conclusions and an outlook to future work.

2 Preliminaries

We introduce here the basic notions of (defeasible) DLs and their inferences. Starting from the two disjoint sets N_C of concept names and N_R of role names, complex concepts can be defined inductively. Let C and D be $\mathcal{EL}$-concepts and $r \in N_R$, then *(complex) $\mathcal{EL}$-concepts* are: *named concepts* ($A \in N_C$), the *top-concept* $\top$, *conjunctions* $C \sqcap D$ and *existential restrictions* $\exists r.C$. The DL $\mathcal{EL}_\bot$ extends $\mathcal{EL}$ by the *bottom-concept* $\bot$, which can be used in conjunctions and existential restrictions. We will occasionally also use the concepts *negation* $\neg C$ and *disjunction* $C \sqcup D$. The semantics of concepts is given by means of interpretations. An *interpretation* $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$ consists of an *interpretation domain* $\Delta^\mathcal{I}$ and a *mapping function* that assigns subsets of the domain $\Delta^\mathcal{I}$ to concept names and binary relations over $\Delta^\mathcal{I}$ to role names. The top-concept is interpreted as the whole domain ($\top^\mathcal{I} = \Delta^\mathcal{I}$) and the bottom-concept as the empty set ($\bot^\mathcal{I} = \emptyset$). The complex concepts are interpreted as follows: conjunction $(C \sqcap D)^\mathcal{I} = C^\mathcal{I} \cap D^\mathcal{I}$, negation $(\neg C)^\mathcal{I} = \Delta^\mathcal{I} \setminus C^\mathcal{I}$, disjunction $(C \sqcup D)^\mathcal{I} = C^\mathcal{I} \cup D^\mathcal{I}$, and existential restriction $(\exists r.C)^\mathcal{I} = \{d \in \Delta^\mathcal{I} \mid \exists e.(d,e) \in r^\mathcal{I} \text{ and } e \in C^\mathcal{I}\}$. If in an interpretation $\mathcal{I}$ $(d,e) \in r^\mathcal{I}$ holds, then e is called a *role successor* of d.

DL ontologies can state (monotonous) relationships between concepts. Let C and D be concepts. A *general concept inclusion axiom* (GCI) is of the form: $C \sqsubseteq D$. A *TBox* $\mathcal{T}$ is a finite set of GCIs. A concept C is *satisfied* by an interpretation $\mathcal{I}$ *iff* $C^\mathcal{I} \neq \emptyset$. A GCI $C \sqsubseteq D$ is *satisfied* in an interpretation $\mathcal{I}$, *iff* $C^\mathcal{I} \subseteq D^\mathcal{I}$ (written $\mathcal{I} \models C \sqsubseteq D$). An interpretation $\mathcal{I}$ is a *model of a TBox* $\mathcal{T}$, *iff* $\mathcal{I}$ satisfies all GCIs in $\mathcal{T}$ (written $\mathcal{I} \models \mathcal{T}$). Based on the notion of a model, DL reasoning problems are defined. A concept is *consistent* w.r.t. a TBox $\mathcal{T}$ *iff* some model of $\mathcal{T}$ satisfies the concept. A concept C is *subsumed by* a concept D w.r.t. a TBox $\mathcal{T}$ (written $C \sqsubseteq_\mathcal{T} D$) *iff* $C^\mathcal{I} \subseteq D^\mathcal{I}$ holds in all models $\mathcal{I}$ of $\mathcal{T}$.

We fix some notation to access parts of knowledge bases or concepts. Let X denote a concept or a TBox, then $sig(X)$ denotes the *signature* of X. We define $sig_{N_C}(X) = sig(X) \cap N_C$ and $sig_{N_R}(X) = sig(X) \cap N_R$. We also define the set

$Qc(X)$ of *quantified concepts* in X as $F \in Qc(X)$ *iff* $\exists r.F$ syntactically occurs in X for some $r \in N_R$. In extensions of $\mathcal{EL}$ that are in the Horn fragment of DLs, canonical models are widely used for reasoning [2]. For an $\mathcal{EL}_\bot$-TBox $\mathcal{T}$, the *canonical model* $\mathcal{I}_\mathcal{T} = (\Delta, \cdot^{\mathcal{I}_\mathcal{T}})$ of $\mathcal{T}$ with $\Delta = \{d_F \mid F \in Qc(\mathcal{T})\}$ has the mapping function satisfying the conditions $d_F \in A^{\mathcal{I}_\mathcal{T}}$ *iff* $F \sqsubseteq_\mathcal{T} A$ and $(d_F, d_G) \in r^{\mathcal{I}_\mathcal{T}}$ *iff* $F \sqsubseteq_\mathcal{T} \exists r.G$. Once the canonical model is computed, subsumption relationships between concepts can be directly read-off from it [2,1].

In defeasible DLs it can be stated that instances of a concept are subsumed by another concept as long as there is no contradicting information. A *defeasible concept inclusion* (DCI) is of the form $C \mathrel{\underset{\sim}{\sqsubset}} D$ and states that C *usually* entails D. A *DBox* $\mathcal{D}$ is a finite set of DCIs. A *defeasible knowledge base* (DKB) $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ consists of a TBox $\mathcal{T}$ and a DBox $\mathcal{D}$. The definitions for $sig(X)$, $sig_{N_C}(X)$, $sig_{N_R}(X)$ and $Qc(X)$ extend to DBoxes or DKBs in the obvious way. A *materialisation* of a DBox $\mathcal{D}$ is the concept $\overline{\mathcal{D}} = \bigsqcap_{E \mathrel{\underset{\sim}{\sqsubset}} F \in \mathcal{D}} (\neg E \sqcup F)$. The semantics of DBoxes differ from the ones for TBoxes, since DCIs need not hold at each element in the model whereas GCIs do. The satisfaction of DCIs for $d \in \Delta^\mathcal{I}$ is captured by $\mathcal{I}, d \models \mathcal{D}$ *iff* $\forall G \mathrel{\underset{\sim}{\sqsubset}} H \in \mathcal{D}. d \in G^\mathcal{I} \implies d \in H^\mathcal{I}$. Usually, the semantics of DBoxes is given by means of ranked/ordered interpretations—called preferential model semantics [3][8]. Instead of using these, we define a new kind of model for DKBs (in Sect. 4) that extends canonical models for $\mathcal{EL}_\bot$. The main idea is to use several copies of the representatives, such as d_F, for each existentially quantified concept, where each copy satisfies a different set of DCIs from the entire lattice built over (the subsets of) the DBox.

We want to develop a decision procedure for (defeasible) subsumption relationships between concepts, say C and D, w.r.t. a given DKB $\mathcal{K}$ under relevant closure. For the remainder of the paper we make two simplifying assumptions for the sake of ease of presentation. We assume w.l.o.g. that (1) concepts C and D appear syntactically in $Qc(\mathcal{T})$ which can be achieved by adding $\exists r.E \sqsubseteq \top$ with $E \in \{C, D\}$ to $\mathcal{T}$ and (2) all quantified concepts in $\mathcal{K}$ are consistent i.e., $\forall F \in Qc(\mathcal{K}). F \not\sqsubseteq_\mathcal{T} \bot$ and thus $\bot \notin Qc(\mathcal{K})$.

To motivate our approach for reasoning under relevant closure in defeasible $\mathcal{EL}_\bot$, we discuss first earlier approaches for this task and identify their main shortcoming.

3 Minimal Relevant Closure by Materialisation

We recall the reduction algorithms for reasoning by Casini et al. from [4]. Rational entailment in [4] uses materialisation of DCIs to decide defeasible subsumptions $C \mathrel{\underset{\sim}{\sqsubset}} D$ w.r.t. a given DKB $\mathcal{K} = (\mathcal{T}, \mathcal{D})$. Since C might be inconsistent w.r.t. the materialisation of the entire DBox $\mathcal{D}$, the algorithm needs to determine a subset $\mathcal{D}' \subseteq \mathcal{D}$ whose materialisation is consistent with C and $\mathcal{T}$ in order to decide whether $\overline{\mathcal{D}'} \sqcap C \sqsubseteq_\mathcal{T} D$ holds. To obtain $\mathcal{D}'$, $\mathcal{D}$ is iteratively reduced to that subset containing all DCIs whose left-hand side is inconsistent in conjunction with the materialisation of the current DBox: $\mathcal{E}(\mathcal{D}) = \{C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D} \mid \mathcal{T} \models \overline{\mathcal{D}} \sqcap C \sqsubseteq \bot\}$. Define $\mathcal{E}^1(\mathcal{D}) = \mathcal{E}(\mathcal{D})$ and $\mathcal{E}^j(\mathcal{D}) = \mathcal{E}(\mathcal{E}^{j-1}(\mathcal{D}))$ (for $j > 1$). Using $\mathcal{E}()$, the DCIs in

$\mathcal{D}$ can be ranked according to their level of exceptionality, i.e., $r_\mathcal{K}(G \sqsubsim H) = i-1$, for the smallest i s.t. $G \sqsubsim H \notin \mathcal{E}^i(\mathcal{D})$, or $r_\mathcal{K}(G \sqsubsim H) = \infty$ if no such i exists. A DKB $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ is *well-separated* if no DCI in $\mathcal{D}$ has an infinite rank of exceptionality [3]. We assume w.l.o.g. that all DKBs are well-separated: any DKB $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ can be transformed into a well-separated DKB $\mathcal{K}'$ by testing a quadratic number of subsumptions in the size of $\mathcal{D}$: $\mathcal{K}' = \big(\mathcal{T} \cup \{C \sqsubseteq \bot \mid r_\mathcal{K}(C \sqsubsim D) = \infty\}, \mathcal{D} \setminus \{C \sqsubsim D \mid r_\mathcal{K}(C \sqsubsim D) = \infty\}\big)$. Based on the level of exceptionality, the algorithm from [4] partitions the DBox $\mathcal{D}$ into $(E_0, E_1, \ldots, E_n)$ where $E_i = \{G \sqsubsim H \in \mathcal{D} \mid r_\mathcal{K}(G \sqsubsim H) = i\}$, i.e. $\mathcal{D} = \bigcup_{i=0}^{n} E_i$. To find the maximal (w.r.t. cardinality) subset $\mathcal{D}'$ of $\mathcal{D}$, whose materialisation is consistent with C and $\mathcal{T}$ the procedure starts with $\mathcal{D}' = \mathcal{D}$. If $\overline{\mathcal{D}'} \sqcap C \sqsubseteq_\mathcal{T} \bot$, then E_i is removed from $\mathcal{D}'$ for the smallest not yet used i. This test and removal is done iteratively until a subset of $\mathcal{D}$ is reached whose materialisation is consistent with C and $\mathcal{T}$.

While rational closure treats inconsistencies with the granularity of the partitions E_i, relevant closure uses a more fine-grained treatment. To illustrate this, let $G \sqsubsim H \in E_0$ and assume that C is only consistent with $\mathcal{D} \setminus E_0$ (or its subsets). It need not hold, that $(\neg G \sqcup H) \sqcap \overline{\mathcal{D} \setminus E_0} \sqcap C \sqsubseteq_\mathcal{T} \bot$, since the inconsistency may be due to other DCIs in E_0. Still $G \sqsubsim H$ is never used for reasoning about C. This effect is called *inheritance blocking*, as it might be possible to include $G \sqsubsim H$ for reasoning about C, but other DCIs induce some inconsistency and so block the inheritance of property $G \sqsubsim H$ for C. Relevant closure disregards only DCIs that are *relevant* for the inconsistency of C, thereby averting inheritance blocking. The notion of relevance is introduced in [4] in terms of justification.

Definition 1. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB, $\mathcal{J} \subseteq \mathcal{D}$ and C a concept. $\mathcal{J}$ is a* C-justification *w.r.t. $\mathcal{K}$ iff $\overline{\mathcal{J}} \sqcap C \sqsubseteq_\mathcal{T} \bot$ and $\overline{\mathcal{J}'} \sqcap C \not\sqsubseteq_\mathcal{T} \bot$ for all $\mathcal{J}' \subset \mathcal{J}$.*

Let *justifications*$(\mathcal{K}, C) = (\mathcal{J}_1, \ldots, \mathcal{J}_m)$ be the (exponential time [9]) procedure that determines and returns all minimal C-justifications w.r.t. $\mathcal{K}$.

To present a simplified (but equivalent) version of the algorithm from [4] for computing minimal relevant closure containment, we need to define the $\mathcal{D}' \subseteq \mathcal{D}$ that is consistent with C and ultimately used for deciding $\overline{\mathcal{D}'} \sqcap C \sqsubseteq_\mathcal{T} D$. Let *partition*$(\mathcal{D}) = (E_0, \ldots, E_n)$ be a function that computes the above defined partitioning of DBoxes and let $\mathcal{J} \subseteq \mathcal{D}$. Then *min*(*partition*$(\mathcal{D}), \mathcal{J})$ returns E_i for the *smallest* i, s.t. $\mathcal{J} \cap E_i \neq \emptyset$. Given $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ and the subsumption query $C \sqsubsim D$, we define the rank-minimal part of all C-justifications w.r.t. $\mathcal{K}$ as $(\mathcal{M}_1, \ldots, \mathcal{M}_m)$ for $(\mathcal{J}_1, \ldots, \mathcal{J}_m) =$ *justifications*$(\mathcal{K}, C)$, where $\mathcal{M}_i = \mathcal{J}_i \cap$ *min*(*partition*$(\mathcal{D}), \mathcal{J}_i)$, for $1 \leq i \leq m$. In order to obtain a subset of $\mathcal{D}$ that is consistent with C, at least one statement from every justification has to be removed from $\mathcal{D}$. By a preference of exceptionality rank and due to a lack of more refined preference[1], the removed statements shall be the rank-minimal parts of all justifications, i.e. $\mathcal{D}' = \mathcal{D} \setminus (\bigcup_{i=1}^{m} \mathcal{M}_i)$. We denote non-monotonic entailments obtained by minimal relevant closure and materialisation as $\mathcal{K} \models_m C \sqsubsim D$ *iff* $\overline{\mathcal{D}'} \sqcap C \sqsubseteq_\mathcal{T} D$ for $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ and $\mathcal{D}'$ as defined above.

[1] Such as a quantitative ranking of DCIs.

The following example illustrates the problem of inheritance blocking caused by rational closure, but not by minimal relevant closure.

Example 2. Let $\mathcal{K}_{ex1} = (\mathcal{T}_{ex1}, \mathcal{D}_{ex1})$ with:

$$\begin{aligned}\mathcal{T}_{ex1} &= \{Boss \sqsubseteq Worker,\ Boss \sqcap \exists superior.Worker \sqsubseteq \bot\},\\ \mathcal{D}_{ex1} &= \{Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Boss,\ Worker \mathrel{\underset{\sim}{\sqsubset}} Productive,\\ &\quad\ Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible\}, \text{ and}\\ partition(\mathcal{D}_{ex1}) &= \big(E_0 = \{Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Boss,\ Worker \mathrel{\underset{\sim}{\sqsubset}} Productive\},\\ &\quad\ E_1 = \{Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible\}\big).\end{aligned}$$

Rational closure recognises the inconsistency $\overline{\mathcal{D}_{ex1}} \sqcap Boss \sqsubseteq_{\mathcal{T}_{ex1}} \bot$, but it holds that $\overline{\mathcal{D}_{ex1} \setminus E_0} \sqcap Boss \not\sqsubseteq_{\mathcal{T}_{ex1}} \bot$. Thus $Boss \mathrel{\underset{\sim}{\sqsubset}} Worker \sqcap Responsible$ is entailed from $\mathcal{K}_{ex1}$, while $Boss \mathrel{\underset{\sim}{\sqsubset}} Productive$ is not, even though the DCI $Worker \mathrel{\underset{\sim}{\sqsubset}} Productive$ does not cause the inconsistency of $Boss$. For minimal relevant closure, $\mathcal{J}_1 = \{Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Boss\}$ is the only $Boss$-justification w.r.t. $\mathcal{K}_{ex1}$. Therefore, the largest consistent DBox subset of $\mathcal{D}_{ex1}$ for reasoning about the concept $Boss$ is $\mathcal{D}' = \{Worker \mathrel{\underset{\sim}{\sqsubset}} Productive,\ Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible\}$, providing the consequence $\overline{\mathcal{D}'} \sqcap Boss \sqsubseteq_{\mathcal{T}_{ex1}} Productive$.

Example 2 also illustrates the issue caused by employing materialisation that is addressed in this paper. Materialising the DCI $Worker \mathrel{\underset{\sim}{\sqsubset}} Productive$ to $\neg Worker \sqcup Productive$ in conjunction with $\exists superior.Worker$ yields a concept that is not subsumed by $\exists superior.Productive$. The defeasible information is unjustly disregarded when reasoning about quantified concepts yielding uniformly non-typical role successors. Hence, in Example 2, both rational and relevant closure (based on materialisation) are oblivious to the conclusion $Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Responsible$.

4 Typicality Models for Propositional Relevant Entailment

In order to achieve relevant entailment also for quantified concepts, DCIs need to hold for concepts in (nested) existential restrictions. A naive idea to extend the materialisation approach is to enrich all concepts in existential restrictions with materialisations of the DBox. However, for Example 2, enriching the concept $\exists superior.Boss$ with $Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Boss$ to $\exists superior.(Boss \sqcap (\neg Worker \sqcup \exists superior.Boss))$ leads to infinitely many such enriching steps (due to $Boss \sqsubseteq Worker$). Instead, our approach is to extend the canonical models for the classical members of the $\mathcal{EL}$-family to DDLs. Our new kind of models captures varying amounts of DCIs from a DKB to be satisfied by role successors. Their domain essentially consist of copies of the domain of a classical canonical model for each set of DCIs. These *typicality models* were recently introduced by us in [12]. To develop the semantics for reasoning under nested relevant entailment and an appropriate reasoning procedure we proceed in two steps:

1. We introduce minimal typicality models with a lattice domain where all domain elements have non-typical role successors only, i.e., no role successor

needs to satisfy any DCI. We show that these minimal typicality models yield exactly the same subsumption relationships as the materialisation-based relevant entailment in [4].

2. We extend minimal typicality models to maximal typicality models, where each role successor required by $\mathcal{K}$ is chosen such that it satisfies a subset of DCIs from $\mathcal{D}$ that is of maximal cardinality while not causing an inconsistency. We define subsumption under nested relevant entailment based on maximal typicality models and show that these models then yield more subsumption relationships than the materialisation-based relevant entailment.

To devise an algorithm that gives the same entailments as materialisation-based relevant entailment, we use the same subsets of the DBox as Casini et al. in [4] based on justifications. To decide the entailment of $C \sqsubsetsim D$ w.r.t. $\mathcal{K} = (\mathcal{T}, \mathcal{D})$, the subset $\mathcal{D}'$ of $\mathcal{D}$ is constructed by removing rank-minimal parts of all justifications relevant for the inconsistency of C. Since we need to distinguish the subsets obtained from C-justifications for different concepts C, we denote from now on, $\mathcal{D}'$ as $\mathcal{D}_C$ (e.g. for $\mathcal{D}_X \subseteq \mathcal{D}$, use X-justifications w.r.t. $\mathcal{K}$).

In order to infer all the undefeated facts for a concept F, the representative domain element of F in a model needs to satisfy the largest subset of $\mathcal{D}$ that is still satisfiable together with the TBox. If minimal relevant closure is used, this subset is obviously $\mathcal{D}_F$. Now, in $\mathcal{EL}_\bot$-concepts a syntactical sub-concept F can occur in multiple existential restrictions (not top-level), causing several role successors in a model. These in turn might be able to satisfy any subset of DCIs "up to" $\mathcal{D}_F$ each. To be able to capture elements satisfying any set of DCIs, typicality interpretations (potentially) need a representative domain element for each subset of the given DBox $\mathcal{D}$. The subsets of $\mathcal{D}$ form a lattice.

Definition 3. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB. A* complete lattice domain *is defined as $\Delta^{\mathcal{K}} = \bigcup_{\mathcal{U} \subseteq \mathcal{D}} \{ d_F^{\mathcal{U}} \mid F \in Qc(\mathcal{K}) \}$. A domain Δ with $\{ d_F^{\emptyset} \mid F \in Qc(\mathcal{K}) \} \subseteq \Delta \subseteq \Delta^{\mathcal{K}}$ is a* lattice domain.

Typicality interpretations over a lattice domain are the basis for our semantics and we define under which conditions a DKB is satisfied in such interpretations.

Definition 4 (model of $\mathcal{K}$). *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB. A typicality interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ over a lattice domain $\Delta^{\mathcal{I}}$ is a model of $\mathcal{K}$ (written $\mathcal{I} \models \mathcal{K}$) iff 1. $\mathcal{I} \models \mathcal{T}$ and 2. $\mathcal{I}, d_F^{\mathcal{U}} \models \mathcal{U}$ for all $\mathcal{U} \subseteq \mathcal{D}$, $F \in Qc(\mathcal{K})$.*

We define when a defeasible subsumption relationship holds in a typicality interpretation over a lattice domain.

Definition 5. *Let $\mathcal{I}$ be a typicality interpretation over a lattice domain. Then $\mathcal{I}$* satisfies a defeasible subsumption *$C \sqsubsetsim D$ (written $\mathcal{I} \models C \sqsubsetsim D$) iff $d_C^{\mathcal{D}_C} \in D^{\mathcal{I}}$.*

In order to construct a model for $\mathcal{K}$ by means of a TBox, we use auxiliary concept names from the set $N_C^{aux} \subseteq N_C \setminus sig(\mathcal{K})$ to introduce representatives for all $F \in Qc(\mathcal{K})$ for each subset of the given DBox. Given concept F and DBox $\mathcal{D}$, we use $F_{\mathcal{D}} \in N_C^{aux}$ to define the *extended TBox* of F:

$$\mathcal{T}_{\mathcal{D}}(F) \;=\; \mathcal{T} \cup \{F_{\mathcal{D}} \sqsubseteq F\} \cup \{F_{\mathcal{D}} \sqcap G \sqsubseteq H \mid G \sqsubsetsim H \in \mathcal{D}\} \qquad (1)$$

Here $\{F_{\mathcal{D}} \sqsubseteq F\}$ ensures that all constraints on F hold for the auxiliary concept as well. The last set of GCIs in Eq. (1) ensures that every element in $F_{\mathcal{D}}^{\mathcal{I}}$ (for $\mathcal{I} \models \mathcal{T}_{\mathcal{D}}(F)$) satisfies the DCIs in $\mathcal{D}$. The auxiliary concept $F_{\emptyset}$ introduced in the extended TBox $\mathcal{T}_{\emptyset}(F)$ and the concept F from $\mathcal{T}$ have the same subsumers.

Proposition 6. *Let $\mathcal{T}$ be a TBox and F, G be concepts with $sig(G) \cap N_C^{aux} = \emptyset$. Then $F \sqsubseteq_{\mathcal{T}} G$ iff $F_{\emptyset} \sqsubseteq_{\mathcal{T}_{\emptyset}(F)} G$.*

The proof for this proposition as well as most results of this paper are given in detail in the technical report [11]. To use typicality interpretations for reasoning under materialisation-based relevant entailment, the DCIs from $\mathcal{U} \subseteq \mathcal{D}$ need to be satisfied at the elements $d_F^{\mathcal{U}}$ representing $F \in Qc(\mathcal{K})$, but not (necessarily) for the role successors of these elements. In fact, it suffices to construct a typicality interpretation of minimally typical role successors (for now), i.e. to use only the TBox for reasoning about role successors induced by existential restrictions.

Definition 7. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB and $\mathcal{U} \subseteq \mathcal{D}$. The* minimal typicality model *over a lattice domain $\mathfrak{L}_{\mathcal{K}}$ of $\mathcal{K}$ consists of the lattice domain $\Delta^{\mathfrak{L}_{\mathcal{K}}} = \{d_F^{\mathcal{U}} \in \Delta^{\mathcal{K}} \mid F_{\mathcal{U}} \not\sqsubseteq_{\mathcal{T}_{\mathcal{U}}(F)} \bot\}$ and an interpretation mapping that satisfies the following conditions for all $d_F^{\mathcal{U}} \in \Delta^{\mathfrak{L}_{\mathcal{K}}}$:*

- *$d_F^{\mathcal{U}} \in A^{\mathfrak{L}_{\mathcal{K}}}$ iff $F_{\mathcal{U}} \sqsubseteq_{\mathcal{T}_{\mathcal{U}}(F)} A$, for $A \in sig_{N_C}(\mathcal{K})$ and*
- *$(d_F^{\mathcal{U}}, d_G^{\emptyset}) \in r^{\mathfrak{L}_{\mathcal{K}}}$ iff $F_{\mathcal{U}} \sqsubseteq_{\mathcal{T}_{\mathcal{U}}(F)} \exists r.G$, for $r \in sig_{N_R}(\mathcal{K})$.*

Typicality models need not use the complete lattice domain of $2^{|\mathcal{D}|} * |Qc(\mathcal{K})|$ elements due to inconsistent combinations of the represented concept F, $\mathcal{U}$ and $\mathcal{T}$. Note that $\mathfrak{L}_{\mathcal{K}}$ is well-defined, as the initial assumption that $F \not\sqsubseteq_{\mathcal{T}} \bot$ holds (for all $F \in Qc(\mathcal{K})$) implies that all $d_F^{\emptyset}$ exist in $\Delta^{\mathfrak{L}_{\mathcal{K}}}$ by Prop. 6. Furthermore, it is not hard to show that Prop. 6 implies that $\mathfrak{L}_{\mathcal{K}}$, restricted to elements regarding the empty set of DCIs, is the classical canonical model for the $\mathcal{EL}_{\bot}$ TBox $\mathcal{T}$.

Example 8 (Minimal typicality model). Consider again the DKB $\mathcal{K}_{ex1}$ from Example 2 with the consistent subsets of the DBox $\mathcal{D}_{Worker} = \mathcal{D}_{ex1}$, and $\mathcal{D}_{Boss} = \{Worker \sqsubsim Productive,\ Boss \sqsubsim Responsible\}$ w.r.t. $Worker$ and $Boss$, respectively. The subset-lattice of $\mathcal{D}_{ex1}$ and $\mathfrak{L}_{\mathcal{K}_{ex1}}$ are illustrated in Fig. 1 using obvious abbreviations. Note, that the domain elements are grouped in *grey boxes* according to the subset-lattice indicating which DBox subsets are satisfied by which domain elements. According to Definition 5, $\mathfrak{L}_{\mathcal{K}_{ex1}} \models Worker \sqsubsim \exists superior.Boss$, as well as $\mathfrak{L}_{\mathcal{K}_{ex1}} \models Boss \sqsubsim Responsible \sqcap Productive$, because $d_{Worker}^{\mathcal{D}_{Worker}}$ and $d_{Boss}^{\mathcal{D}_{Boss}}$ satisfy $\mathcal{D}_{Worker}$ and $\mathcal{D}_{Boss}$, respectively. Rational closure would select the DBox subset $\{Boss \sqsubsim Responsible\}$ for reasoning about $Boss$ and would therefore not conclude the latter subsumption.

For the minimal typicality model $\mathfrak{L}_{\mathcal{K}}$ it can be shown that $d_C^{\mathcal{D}_C} \in D^{\mathfrak{L}_{\mathcal{K}}}$ (Def. 5) is equivalent to reasoning with the extended TBox, i.e. deciding $C_{\mathcal{D}_C} \sqsubseteq_{\mathcal{T}_{\mathcal{D}_C}(C)} D$. The following claim can be proved using this equivalence.

Lemma 9. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB. Then the minimal typicality model $\mathfrak{L}_{\mathcal{K}}$ over the lattice domain $\Delta^{\mathfrak{L}_{\mathcal{K}}}$ is a model of $\mathcal{K}$.*

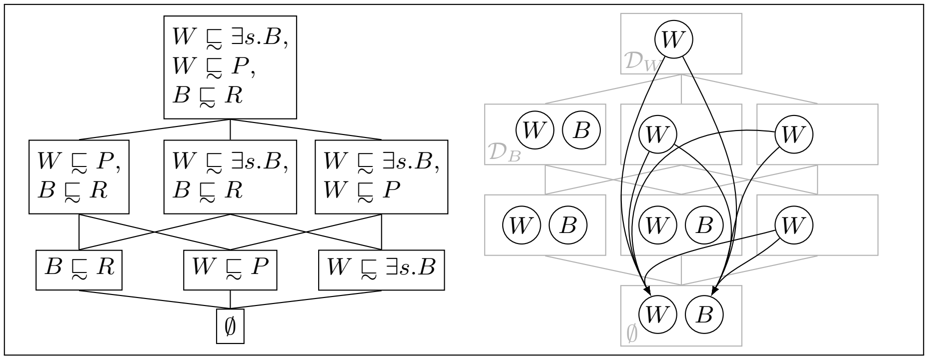

Fig. 1: a Subset lattice of $\mathcal{D}_{ex1}$ and **b** $\mathfrak{L}_{\mathcal{K}_{ex1}}$

We want to characterise different entailment relations based on different kinds of typicality models for a given DKB $\mathcal{K}$ which vary in the defeasible information admitted for required role successors. We use minimal typicality models over a lattice domain to characterise entailment of propositional nature $\models_p$.

Definition 10. *Let $\mathcal{K}$ be a DKB. $\mathcal{K}$* propositionally entails *a defeasible subsumption relationship $C \sqsubsetsim D$ (written $\mathcal{K} \models_p C \sqsubsetsim D$) iff $\mathfrak{L}_\mathcal{K} \models C \sqsubsetsim D$.*

This form of entailment is called *propositional* since all role successors are uniformly non-typical and thus DCIs are neglected for quantified concepts. Next, we investigate the relationship between $\models_m$ (Sec. 3) and $\models_p$. Our approach to decide propositional entailments based on the extended TBox for a concept F, coincides with enriching F with the materialisation of the given DBox.

Lemma 11. *Let $\mathcal{T}$ be a TBox $\mathcal{T}$, $\mathcal{D}$ a DBox, and C, D be concepts, with $sig(X) \cap N_C^{aux} = \emptyset$ (for $X \in \{\mathcal{T}, \mathcal{D}, C, D\}$). Then $\overline{\mathcal{D}} \sqcap C \sqsubseteq_\mathcal{T} D$ iff $C_\mathcal{D} \sqsubseteq_{\mathcal{T}_\mathcal{D}(C)} D$.*

Proof (sketch). The lemma is proven by induction on the size of $\mathcal{D}$. The base case is $\mathcal{D} = \emptyset$ and thereby Prop. 6 holds. For the induction step, let $\mathcal{D}' = \mathcal{D} \cup \{G \sqsubsetsim H\}$ and use the hypothesis that the claim holds for $\mathcal{D}$. We do a case distinction for (i) $\overline{\mathcal{D}} \sqcap C \sqsubseteq G$ and (ii) $\overline{\mathcal{D}} \sqcap C \not\sqsubseteq G$. For case (i), it can be shown that reasoning with $C \sqcap H$, yields the same consequences as reasoning with C. All elements in $C \sqcap H$ satisfying G, already satisfy H. Thus we can remove the introduced DCI $G \sqsubsetsim H$ without losing any consequences. Hence, $G \sqsubsetsim H$ can be removed from $\mathcal{D}'$ and the induction hypothesis holds. In case (ii) we show that the added DCI has no effect on the reasoning. By the condition of this case, no element in C, satisfying all DCIs in D, satisfies G. Therefore, the presence of the DCI $G \sqsubsetsim H$ does not affect the reasoning allowing it to be removed in order to reduce the induction step to the induction hypothesis. In both cases, both sides of the "*iff*" in the claim are reduced to their respective sides of the induction hypothesis individually. Since the full proof is fairly long and technical, it is deferred to the technical report [11].

Although the entailment relations $\models_m$ as introduced in [4] and $\models_p$ are defined in different ways and are based on distinct semantics, they yield the same consequences (for subsumption) w.r.t. DKBs.

Theorem 12. *$\mathcal{K} \models_p C \mathrel{\underset{\sim}{\sqsubset}} D$ iff $\mathcal{K} \models_m C \mathrel{\underset{\sim}{\sqsubset}} D$.*

Proof (sketch). $\mathcal{K} \models_p C \mathrel{\underset{\sim}{\sqsubset}} D$ is defined as $\mathfrak{L}_\mathcal{K} \models C \mathrel{\underset{\sim}{\sqsubset}} D$, i.e. $d_C^{\mathcal{D}_C} \in D^{\mathfrak{L}_\mathcal{K}}$ which is equivalent to $C_{\mathcal{D}_C} \sqsubseteq_{\mathcal{T}_{\mathcal{D}_C}(C)} D$ as shown in the technical report [11]. This subsumption in turn is equivalent to deciding $\overline{\mathcal{D}_C} \sqcap C \sqsubseteq_\mathcal{T} D$ by Lemma 11, which is precisely the definition of $\mathcal{K} \models_m C \mathrel{\underset{\sim}{\sqsubset}} D$ (from Section 3). □

In addition, this result shows that entailments based on minimal typicality models also bear the shortcomings for defeasible reasoning regarding nested existential restrictions—a nuisance which we want to alleviate next.

5 Maximal Typicality Models for Relevant Entailment

We illustrate by continuing on Example 8 how defeasible information is disregarded for nested existential restrictions and our proposed countermeasure.

Example 13. Consider again the DKB $\mathcal{K}_{ex1}$ from Example 8, with $\mathfrak{L}_{\mathcal{K}_{ex1}}$ (as depicted in Fig. 1). No defeasible information is used for reasoning over the *superior* successors of $d_{Worker}^{\mathcal{D}_{Worker}}$ and thus $\mathfrak{L}_\mathcal{K} \not\models Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Responsible$. However, $Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible$ remains *undefeated* for $d_{Boss}^{\emptyset}$. Instead of satisfying $Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible$ at $d_{Boss}^{\emptyset}$, we can "upgrade" the existing *superior* relationship, e.g. $(d_{Worker}^{\mathcal{D}_{Worker}}, d_{Boss}^{\emptyset})$ to $(d_{Worker}^{\mathcal{D}_{Worker}}, d_{Boss}^{\{Boss \mathrel{\underset{\sim}{\sqsubset}} Responsible\}})$ —as illustrated in Fig. 2 by the *dashed arrows*. Note that Figure 2 depicts only an excerpt of $\mathfrak{L}_{\mathcal{K}_{ex1}}$ for comprehensibility. Our method upgrades typicality of role successors as much as possible, as long as it does not result in inconsistencies. Here, this even yields the conclusion $Worker \mathrel{\underset{\sim}{\sqsubset}} \exists superior.Productive$.

Upgrading the typicality of a role successor depends on the information present in the model. Different orders of such upgrade steps can yield different models of increased typicality. In order to handle sets of models over the same lattice domain Δ, we need the notions of intersection and inclusion of models.

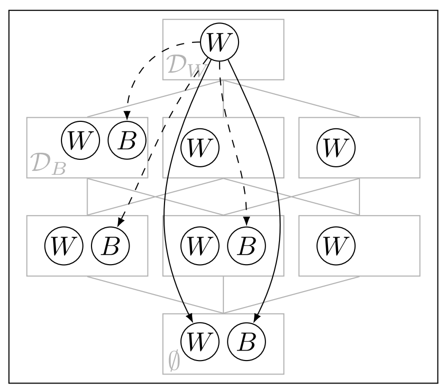

Fig. 2: Increasing typicality of role successors

Definition 14. *For two interpretations $\mathcal{I}, \mathcal{J}$ over the same domain Δ we define*

- $\mathcal{I} \cap \mathcal{J} = (\Delta, \cdot^{\mathcal{I}\cap\mathcal{J}})$ *with* $A^{\mathcal{I}\cap\mathcal{J}} = A^\mathcal{I} \cap A^\mathcal{J}$ *(for $A \in N_C$) and* $r^{\mathcal{I}\cap\mathcal{J}} = r^\mathcal{I} \cap r^\mathcal{J}$ *(for $r \in N_R$)*
- $\mathcal{I} \subseteq \mathcal{J}$ *iff* $\forall A \in N_C.A^\mathcal{I} \subseteq A^\mathcal{J}$ *and* $\forall r \in N_R.r^\mathcal{I} \subseteq r^\mathcal{J}$.

Using this definition we can formalise the notion of "upgrading the typicality" of a typicality interpretation, i.e. introduce copies of role edges that point to elements representing the same concept, but satisfying more defeasible statements.

Definition 15. *Let $\mathcal{I}$ be a typicality interpretation over a lattice domain for $\mathcal{K} = (\mathcal{T}, \mathcal{D})$. The* set of more typical role edges *for a given role r is defined as*

$$TR_{\mathcal{I}}(r) = \{(d_G^{\mathcal{X}}, d_H^{\mathcal{U}'}) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \setminus r^{\mathcal{I}} \mid \exists \mathcal{U} \subseteq \mathcal{D}.(d_G^{\mathcal{X}}, d_H^{\mathcal{U}}) \in r^{\mathcal{I}} \wedge \mathcal{U} \subset \mathcal{U}' \subseteq \mathcal{D}_H\}.$$

Let $\mathcal{I}$ and $\mathcal{J}$ be typicality interpretations. $\mathcal{J}$ is a typicality extension *of $\mathcal{I}$ iff 1. $\Delta^{\mathcal{J}} = \Delta^{\mathcal{I}}$, 2. $A^{\mathcal{J}} = A^{\mathcal{I}}$ (for $A \in N_C$), 3. $r^{\mathcal{J}} = r^{\mathcal{I}} \cup R$, where $R \subseteq TR_{\mathcal{I}}(r)$ (for $r \in sig_{N_R}(\mathcal{K})$), and 4. $\exists r \in sig_{N_R}(\mathcal{K})$. $r^{\mathcal{I}} \subset r^{\mathcal{J}}$. The* set of all typicality extensions *of a typicality interpretation $\mathcal{I}$ is $typ(\mathcal{I})$.*

With typicality extensions at hand we can transform typicality interpretations into a set of more typical interpretations. Unfortunately, this operation does not preserve the property of being a typicality model. Let us demonstrate this by Example 13, let $\mathcal{K}_{ex2} = (\mathcal{T}_{ex2}, \mathcal{D}_{ex1})$, and $\mathcal{T}_{ex2} = \mathcal{T}_{ex1} \cup \{\exists superior.Responsible \sqsubseteq \exists coworker.Worker\}$. Since $\mathfrak{L}_{\mathcal{K}_{ex2}}$ coincides with $\mathfrak{L}_{\mathcal{K}_{ex1}}$, Fig. 2 depicts a typicality extension of $\mathfrak{L}_{\mathcal{K}_{ex2}}$ according to Def. 15. However, the extension in Fig. 2 is no longer a model of $\mathcal{T}_{ex2}$, as the newly introduced GCI is no longer satisfied for $d^{\mathcal{D}}_{Worker}$. It can be extended to a model by introducing a *coworker* successor for $d^{\mathcal{D}}_{Worker}$ that belongs to $Worker$. In order not to introduce unwanted inconsistencies, the successor in this new relationship needs to be picked such that it only contains the information *strictly* required by $\mathcal{K}$, i.e. $d^{\emptyset}_{Worker}$. We formalise the particular *model completions* that we are interested in.

Definition 16. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ be a DKB and Δ a lattice domain. An interpretation $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ is a* model completion *of an interpretation $\mathcal{J} = (\Delta, \cdot^{\mathcal{J}})$ iff 1. $\mathcal{J} \subseteq \mathcal{I}$, 2. $\mathcal{I} \models \mathcal{K}$, and 3. $\forall E \in Qc(\mathcal{K}).d_F^{\mathcal{U}} \in (\exists r.E)^{\mathcal{I}} \implies (d_F^{\mathcal{U}}, d_E^{\emptyset}) \in r^{\mathcal{I}}$ (for any $F \in Qc(\mathcal{K})$ and $\mathcal{U} \subseteq \mathcal{D}$). The set of all model completions of $\mathcal{J}$ is denoted as $mc(\mathcal{J})$.*

An interpretation that is a model completion to itself is called a *safe model* and obviously satisfies the properties of Def. 16, e.g. for some typicality interpretation $\mathcal{J}$ over a lattice domain, all interpretations in $mc(\mathcal{J})$ are safe models.

Since model completions introduce minimal typical role successors they may necessitate further typicality extensions. So, typicality extensions and model completions need to be applied alternatingly until a maximum is reached. Maximality for typicality extensions is characterised in the following way: a typicality interpretation $\mathcal{I}$ is *typicality extensible iff* $\exists \mathcal{J} \in typ(\mathcal{I}).mc(\mathcal{J}) \neq \emptyset$. Intuitively, a typicality interpretation is typicality extensible if it admits to some typicality extension that is, or can be completed to a safe model. Therefore, a typicality interpretation is *maximal iff* it is not typicality extensible. To formalise the process of increasing typicality and completing to a model until reaching maximal typicality, we introduce some notation and an upgrade operator. Given a lattice domain Δ, define the *set of all safe models* over Δ as $P(\Delta) = \{\mathcal{J} \mid \mathcal{J} = (\Delta, \cdot^{\mathcal{J}}) \wedge \mathcal{J} \in mc(\mathcal{J})\}$.

Definition 17. *The typicality upgrade operator $T : 2^{P(\Delta)} \to 2^{P(\Delta)}$ is defined for $S \subseteq P(\Delta)$ as:*

1. $T(S) = S \setminus \{\mathcal{I}\} \cup \bigcup_{\mathcal{J}\in typ(\mathcal{I})} mc(\mathcal{J})$, *if* $\mathcal{I} \in S$ *is typicality extensible,*
2. $T(S) = S$, *otherwise.*

For a given set of model completions $S \subseteq P(\Delta)$*, the* fixpoint of T *is* $T_m(S)$ *if* $T_m(S) = T_{m+1}(S)$ *with* $T_0(S) = S$, $T_i(S) = T(T_{i-1}(S))$ $(i > 0)$*. The* set of maximal typicality extensions *of the typicality interpretations in* S *is* $typ^{max}(S) = T_m(S)$.

It is clear that neither Definition 15 nor 16 supply unique typicality extensions or model completions, respectively. Thus the typicality upgrade operator is defined over sets of typicality interpretations. Applying T to $\{\mathcal{I}\}$ easily leads to an exponential amount of typicality interpretations. The following example illustrates that this typicality extension can quickly lead to multiple different maximal typicality interpretations, starting from a single interpretation.

Example 18. We extend the DKB from Example 8 to DKB $\mathcal{K}_{ex3} = (\mathcal{T}_{ex3}, \mathcal{D}_{ex1})$ with the TBox $\mathcal{T}_{ex3} = \mathcal{T}_{ex1} \cup \{\exists superior.\exists superior.Responsible \sqsubseteq \bot\}$. Let the role edge $(d^{\mathcal{D}}_{Worker}, d^{\emptyset}_{Worker}) \in superior^{\mathfrak{L}_{\mathcal{K}_{ex3}}}$ be upgraded to $(d^{\mathcal{D}}_{Worker}, d^{\mathcal{D}}_{Worker})$ and likewise $(d^{\mathcal{D}}_{Worker}, d^{\emptyset}_{Boss}) \in superior^{\mathfrak{L}_{\mathcal{K}_{ex3}}}$ to $(d^{\mathcal{D}}_{Worker}, d^{\mathcal{D}_{Boss}}_{Boss})$. If both of these upgrades exist in the same typicality extension $\mathcal{J}$, it does not admit to a model completion, as $d^{\mathcal{D}}_{Worker} \in (\exists superior.\exists superior.Responsible)^{\mathcal{J}}$. The typicality upgrade $(d^{\mathcal{D}}_{Worker}, d^{\{Worker \sqsubsim Productive\}}_{Boss})$, however, is "allowed" to occur in a typicality extension, leading to the entailment of $Worker \sqsubsim \exists superior.(Boss \sqcap Productive)$. This shows that inheritance blocking can be remedied even for quantified concepts when upgrading typicality of successors in a lattice domain.

It is clear that the above described process leads to a variety of maximal typicality models, and there are several ways to use these sets for our new entailment. Since in classical DL reasoning entailment considers all models, we pick semantics closely related to *cautious* reasoning here. To this end we build a single model that is canonical in the sense that it is contained in all maximal typicality models obtained from $\mathfrak{L}_{\mathcal{K}}$.

Definition 19. *The* relevant canonical model *is* $\mathfrak{RE}_{\mathcal{K}} = \bigcap_{\mathcal{I}\in typ^{max}(\{\mathfrak{L}_{\mathcal{K}}\})} \mathcal{I}$.

Note that the intersection over all maximal typicality models is well-defined as $typ^{max}(\{\mathfrak{L}_{\mathcal{K}}\})$ is finite and since $\mathfrak{L}_{\mathcal{K}}$ is a safe model (i.e. $\mathfrak{L}_{\mathcal{K}} \in mc(\mathfrak{L}_{\mathcal{K}})$) it is not empty, as shown in the technical report [11].

Lemma 20. *The relevant canonical model* $\mathfrak{RE}_{\mathcal{K}}$ *is a model of the DKB* $\mathcal{K}$.

The main argument in the proof of Lemma 20 is that the intersection of models in any set $S \subseteq P(\Delta)$ yields another model in $P(\Delta)$. This result is ensured by Condition 3 of Definition 16 The relevant canonical model is used to decide nested relevant entailment of the form $C \sqsubsim_{\mathcal{K}} D$, which requires to propagate DCIs to concepts occurring in existential restrictions. We capture this stronger and quantifier-aware relevant entailment.

Definition 21. *Let* $\mathcal{K}$ *be a DKB. A defeasible subsumption relationship* $C \sqsubsim D$ *holds under* nested relevant entailment *(written* $\mathcal{K} \models_q C \sqsubsim D$*) iff* $\mathfrak{RE}_{\mathcal{K}} \models C \sqsubsim D$.

We are ready to state our main result: nested relevant entailment allows for strictly more inferences than the materialisation-based relevant entailment from [4] to compute the relevant closure.

Theorem 22. *For two $\mathcal{EL}_\perp$ concepts C, D and an $\mathcal{EL}_\perp$ DKB $\mathcal{K}$*

1. $\mathcal{K} \models_m C \sqsubsetsim D \implies \mathcal{K} \models_q C \sqsubsetsim D$, *and*
2. $\mathcal{K} \models_m C \sqsubsetsim D \not\Longleftarrow \mathcal{K} \models_q C \sqsubsetsim D$

The first Claim of Theorem 22 is not too hard to show, as $\mathcal{J} \in typ^{max}(\{\mathfrak{L}_\mathcal{K}\})$ implies that $\mathfrak{L}_\mathcal{K} \subseteq \mathcal{J}$ and thus $\mathfrak{L}_\mathcal{K} \subseteq \mathfrak{RE}_\mathcal{K}$. Hence by Definition 5, $\mathfrak{L}_\mathcal{K} \models C \sqsubsetsim D$ implies $\mathfrak{RE}_\mathcal{K} \models C \sqsubsetsim D$. The second claim can be supported by Example 13 and is formally shown in the technical report [11].

Extending our work on rational closure [12], the present approach fixes a deficit induced by materialisation and computes strictly more entailments than materialisation-based relevant reasoning. Since relevant closure is stronger than rational closure, the reduction presented here preserves all merits for rational closure from [12] and fixes remaining problems such as inheritance blocking.

6 Conclusions and Future Work

In this paper we have extended a new approach for reasoning in DDLs to characterise entailment under relevant closure (for deciding subsumption) in the DDL $\mathcal{EL}_\perp$. The new approach is motivated by the observation that earlier reasoning procedures for this problem do not treat existential restrictions in an adequate way. The key idea is to extend canonical models such that for each concept from the DKB, several copies representing different amounts of defeasible information are introduced. The new semantics supports the propagation of defeasible information to concepts nested in existential restrictions. It allows for extensions to more expressive approaches, e.g. more expressive DLs. While rational closure needs to consider only one sequence of increasing subsets of the DBox [12], relevant closure needs (potentially) *all* subsets of the given defeasible information—forming a lattice. In minimal typical models over a lattice domain the role successors are "a-typical" in the sense that they satisfy only the GCIs from the TBox. Such models can be computed by a reduction to classical TBox reasoning. We showed that the obtained entailments coincide with the ones obtained by earlier materialisation-based algorithms. We extended these models to maximally typical models, which have role successors of "maximal typicality". Entailment over these models propagates defeasible information to role successors and thus allows for more entailments.

There are several paths for future work. Besides the extensions to more expressive DLs, an extension to ABox reasoning, i.e., reasoning about data, would be an interesting topic to investigate. Furthermore, a completion-like algorithm as for classical $\mathcal{EL}$ [2,1] would be desirable to effectively compute these models. The current definition of typicality extensions and model completions leaves plenty of room for developing practical algorithms worth implementing.

References

1. Baader, F., Brandt, S., Lutz, C.: Pushing the $\mathcal{EL}$ envelope. In: Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence IJCAI-05. Edinburgh, UK (2005)
2. Brandt, S.: Polynomial time reasoning in a description logic with existential restrictions, GCI axioms, and—what else? In: R.L. de Mantáras, L. Saitta (eds.) Proceedings of the 16th European Conference on Artificial Intelligence (ECAI-2004), pp. 298–302 (2004)
3. Britz, A., Casini, G., Meyer, T., Moodley, K., Varzinczak, I.: Ordered Interpretations and Entailment for Defeasible Description Logics. Tech. rep., CAIR, CSIR Meraka and UKZN, South Africa (2013)
4. Casini, G., Meyer, T., Moodley, K., Nortjé, R.: Relevant closure: A new form of defeasible reasoning for description logics. In: Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA'14), pp. 92–106 (2014)
5. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In: Proceedings of 12th European Conference on Logics in Artificial Intelligence (Jelia'10), pp. 77–90 (2010)
6. Casini, G., Straccia, U.: Defeasible inheritance-based description logics. In: IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, pp. 813–818 (2011)
7. Casini, G., Straccia, U.: Lexicographic closure for defeasible description logics. In: Proceedings of the Eighth Australasian Ontology Workshop, pp. 28–39 (2012)
8. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Semantic characterization of rational closure: From propositional logic to description logics. Artif. Intell. **226**, 1–33 (2015)
9. Horridge, M.: Justification based explanation in ontologies. Ph.D. thesis, University of Manchester (2011)
10. Kraus, S., Lehmann, D.J., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artif. Intell. **44**(1-2), 167–207 (1990)
11. Pensel, M., Turhan, A.Y.: Including quantification in defeasible reasoning for the description logic $\mathcal{EL}_{\bot}$. LTCS-Report 17-01, Chair for Automata Theory, Institute for Theoretical Computer Science, Technische Universität Dresden (2017). See http://lat.inf.tu-dresden.de/research/reports.html
12. Pensel, M., Turhan, A.Y.: Including quantification in defeasible reasoning for the description logic $\mathcal{EL}_{bot}$. In: M. Balduccini, T. Janhunen (eds.) Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning - LPNMR, pp. 78–84. Springer (2017). DOI 10.1007/978-3-319-61660-5_9