$\mathcal{ALC} + \mathrm{T_R^{exp}}$: beyond most likely scenarios in preferential Description Logics of typicality

Gian Luca Pozzato[1]

Dip. Informatica - Università di Torino - Italy
gianluca.pozzato@unito.it

Abstract. In this work we continue our investigation about the opportunity of reasoning about alternative, surprising scenarios in preferential Description Logics of typicality. We extend the results provided in [1] and presented at CILC 2015, where a nonmonotonic procedure for preferential Description Logics has been outlined in order to solve a problem coming from sports entertainment. Here we provide complexity results for the Description Logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$, extending the Description Logic $\mathcal{ALC} + \mathbf{T_R}$ by allowing inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$, where d is a degree of expectedness, and also allowing to reason about surprising extensions of an ABox satisfying cardinality restrictions on concepts. Here we propose a decision procedure for reasoning in $\mathcal{ALC} + \mathbf{T_R^{exp}}$, and we exploit it to show that entailment is in EXPTIME, allowing us to conclude that reasoning in $\mathcal{ALC} + \mathbf{T_R^{exp}}$ is essentially inexpensive.

1 Introduction

Nonmonotonic extensions of Description Logics (from now on, DLs for short) have been actively investigated since the early 90s [2–9] in order to tackle the problem of representing *prototypical* properties of classes and to reason about *defeasible* inheritance. A simple but powerful nonmonotonic extension of DLs is proposed in [10–14]: in this approach "typical" or "normal" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying DL. The semantics of the $\mathbf{T}$ operator is characterized by the core properties of nonmonotonic reasoning axiomatized by either *preferential logic* [15] or *rational logic* [16]. We focus on the Description Logic $\mathcal{ALC} + \mathbf{T_R}$ introduced in [14]. In this logic one can express defeasible inclusions such as "normally, depressed people have sleep disorders":

$$\mathbf{T}(Depressed) \sqsubseteq \exists HasSympthom.SleepDisorder$$

As a difference with standard DLs, one can consistently express exceptions and reason about defeasible inheritance as well. For instance, a knowledge base can consistently express that "normally, a patient affected by depression is not able to react to positive life events", whereas "mood reactivity (ability to feel better temporarily in response to positive life events) is a typical symptom of atypical depression":

$AtypicalDepressed \sqsubseteq Depressed$
$\mathbf{T}(Depressed) \sqsubseteq \neg\exists HasSympthom.MoodReactivity$
$\mathbf{T}(AtypicalDepressed) \sqsubseteq \exists HasSympthom.MoodReactivity$

From a semantic point of view, models of $\mathcal{ALC} + \mathbf{T_R}$ are standard models extended by a function f which selects the typical/most normal instances of any concept C, i.e. the extension of $\mathbf{T}(C)$ is defined as $(\mathbf{T}(C))^{\mathcal{I}} = f(C^{\mathcal{I}})$. The function f satisfies a set of postulates that are a restatement of Kraus, Lehmann, and Magidor's axioms of rational logic $\mathbf{R}$. This allows the typicality operator to inherit well-established properties of nonmonotonic reasoning.

The logic $\mathcal{ALC} + \mathbf{T_R}$ results to be too weak in several application domains. Indeed, although the operator $\mathbf{T}$ is nonmonotonic ($\mathbf{T}(C) \sqsubseteq E$ does not imply $\mathbf{T}(C \sqcap D) \sqsubseteq E$), the logic $\mathcal{ALC} + \mathbf{T_R}$ is monotonic, in the sense that if the fact F follows from a given knowledge base KB, then F also follows from any KB' $\supseteq$ KB. As a consequence, unless a KB contains explicit assumptions about typicality of individuals, there is no way of inferring defeasible properties about them: in the above example, if KB contains the fact that Kate is a depressed woman, i.e. $Depressed(kate)$ belongs to KB, it is not possible to infer that she has sleep disorders ($\exists HasSympthom.SleepDisorder(kate)$). This would be possible only if the KB contained the stronger information that Kate is a *typical* depressed woman, i.e. $\mathbf{T}(Depressed)(kate)$ belongs to (or can be inferred from) KB. In order to overwhelm this limit and perform useful inferences, in [14] the authors have introduced a nonmonotonic extension of the logic $\mathcal{ALC} + \mathbf{T_R}$ based on a minimal model semantics, corresponding to a notion of *rational closure* as defined in [16] for propositional logic. Intuitively, the idea is to restrict our consideration to (canonical) models that maximize typical instances of a concept when consistent with the knowledge base. The resulting logic, call it $\mathcal{ALC} + \mathbf{T_R}^{RaCl}$, supports typicality assumptions, so that if one knows that Kate is depressed, one can nonmonotonically assume that she is also a *typical* depressed if this is consistent, and therefore that she has sleep disorders. From a semantic point of view, the logic $\mathcal{ALC} + \mathbf{T_R}^{RaCl}$ is based on a preference relation among $\mathcal{ALC} + \mathbf{T_R}$ models and a notion of *minimal entailment* restricted to models that are minimal with respect to such preference relation.

The logic $\mathcal{ALC} + \mathbf{T_R}^{RaCl}$ imposes to consider *all* consistent typicality assumptions that are consistent with a given KB. This seems to be too strong in several application domains, in particular when the need arises of bounding the cardinality of the extension of a given concept, that is to say the number of domain elements being members of such a concept, as introduced in [17]. As a further example, consider the following KB from the domain of sports entertainment taken from [1]: $\mathbf{T}(FaceWrestler) \sqsubseteq RoyalRumbleWinner$, $\mathbf{T}(Returning) \sqsubseteq RoyalRumbleWinner$, $\mathbf{T}(Predicted) \sqsubseteq RoyalRumbleWinner$. The first inclusion represents that, normally, a face wrestler wins the Royal Rumble match, an annual wrestling event involving thirty athletes. The second one states that, typically, an athlete returning from an injury wins the Royal Rumble match. The third and last inclusion represents that an athlete whose victory has been predicted by wrestling web sites normally wins the Royal Rumble match. If the assertional part of the KB contains the facts: $FaceWrestler(dean)$, $Returning(dave)$, $FaceWrestler(roman)$, $Predicted(roman)$, whose meaning is that Dean is a face athlete, Dave is returning from an injury, and that Roman is a face wrestler who has been predicted to win the Royal Rumble match, respectively, then in $\mathcal{ALC} + \mathbf{T_R}^{RaCl}$ we conclude that: $\mathbf{T}(FaceWrestler)(dean)$, $\mathbf{T}(Returning)(dave)$, $\mathbf{T}(FaceWrestler)(roman)$, $\mathbf{T}(Predicted)(roman)$, and then that Dean, Dave and Roman are all winners. This hap-

pens in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ because it is consistent to make the three assumptions above, that hold in all minimal models, however one should be interested in three distinct scenarios that cannot be captured by $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ as it is. One could think of extending the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ by means of *cardinality restrictions*, in the example by imposing that there is only one member of the extension of the concept $RoyalRumbleWinner$, however the resulting knowledge base would be inconsistent.

Furthermore, it is sometimes useful to restrict reasoning to *surprising* scenarios, excluding "trivial"/"obvious" ones. For instance, recently a great attention has been devoted to *serendipitous* search engines, that must be able to provide results that are "*surprising, semantically cohesive*, i.e. *relevant* to some information need of the user, or just *interesting*" [18]. In this sense, the scenario (among those satisfying cardinality restrictions) obtained by assuming the largest set of consistent typicality assumptions in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ corresponds to the most trivial one, whereas one could be interested in less expected ones, in which some typicality assumptions are discarded.

In [1] we have moved a first step towards the definition of an extension of Description Logics of typicality allowing to reason about surprising scenarios in presence of cardinality restrictions. In that work, an extension of DL-Lite$_{core}$ has been introduced in order to tackle a problem coming from sports entertainment, namely to find a plausible but surprising outcome of a wrestling event. However, neither decision procedures nor complexity results are provided, being that work only a preliminary contribution in this field. Here we try to move a further step by extending our approach to the more expressive Description Logic $\mathcal{ALC}$. Moreover, whereas the approach in [1] is based on a nonmonotonic extension of DLs based on preferential logic $\mathbf{P}$ [15], here we exploit the nonmonotonic extension of $\mathcal{ALC}$ called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, based on rational models [16] and corresponding to a notion of *rational closure* for Description Logics introduced in [14]. The approach we propose in this work is based on the combination of two components. On the one hand, we allow to express different *degrees of expectedness* of typicality inclusions: this allows to describe several plausible scenarios by considering different combinations of typicality assumptions about individuals named in the ABox. Such degrees introduce a rank of expectedness among plausible scenarios, ranging from surprising to obvious ones. On the other hand, TBoxes are extended to allow restrictions about the cardinality of concepts, in order to "filter" such plausible scenarios. Finally, reasoning tasks are restricted to reasonable but "surprising enough" (or "not obvious") scenarios satisfying cardinality restrictions. In detail, the original contribution of this work can be summarized as follows:

- we introduce a new Description Logic of typicality, called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$, allowing to express a degree of expectedness of typicality inclusions and cardinality restrictions on concepts: TBoxes are extended by (i) inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$ where d is a positive integer, such that an inclusion with degree d is more "trivial" (or "obvious") with respect to another one with degree $d' \leq d$, as well as by (ii) restrictions on the cardinality of concepts of the form $(\odot \, n \, C)$, where $\odot \in \{\leq, \geq, =\}$ and $n \in \mathbb{N}^+$;

- we introduce a notion of extension of an ABox for the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$, corresponding to a set of typicality assumptions that can be performed in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ for individual constants, then we introduce an order relation among extensions whose basic idea is to prefer extensions representing more surprising scenarios;

- we define notions of skeptical and credulous entailment in $\mathcal{ALC} + \mathbf{T_R^{exp}}$, relying on reasoning in $\mathcal{ALC} + \mathbf{T_R}$, but allowing to restrict our concern to "non trivial" scenarios, corresponding to minimal extensions with respect to the order relation among extensions of the previous point;

- we describe a procedure for reasoning in $\mathcal{ALC} + \mathbf{T_R^{exp}}$, that can be used to estimate the complexity of entailment (EXPTIME for both skeptical and credulous entailment).

It is worth noticing that the proposed logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$ is not intended to replace existing extensions of DLs for representing and reasoning about prototypical properties and defeasible inheritance. The idea is that, in some applications, the need of reasoning about surprising scenarios could help domain experts to achieve their goals, wherever standard reasoning is not enough to do it: as an example, in medical diagnosis, the most likely explanation for a set of symptoms is not always the solution to the problem, whereas reasoning about surprising scenarios could help the medical staff in taking alternative explanations into account. In other words, the Description Logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$ is not intended to replace existing nonmonotonic DLs, but to tile them in order to reason about alternative, plausible scenarios when it is needed to go beyond most likely solutions.

2 Description Logics of Typicality

2.1 The monotonic logic $\mathcal{ALC} + \mathbf{T_R}$

The logic $\mathcal{ALC} + \mathbf{T_R}$ is obtained by adding to standard $\mathcal{ALC}$ the typicality operator $\mathbf{T}$ [14]. The intuitive idea is that $\mathbf{T}(C)$ selects the *typical* instances of a concept C. We can therefore distinguish between the properties that hold for all instances of concept C ($C \sqsubseteq D$), and those that only hold for the normal or typical instances of C ($\mathbf{T}(C) \sqsubseteq D$).

The semantics of the $\mathbf{T}$ operator can be formulated in terms of *rational models*: a model $\mathcal{M}$ is any structure $\langle \Delta^\mathcal{I}, <, .^\mathcal{I} \rangle$ where $\Delta^\mathcal{I}$ is the domain, $<$ is an irreflexive, transitive, well-founded and modular (for all x, y, z in Δ, if $x < y$ then either $x < z$ or $z < y$) relation over Δ. In this respect, $x < y$ means that x is "more normal" than y, and that the typical members of a concept C are the minimal elements of C with respect to this relation. An element $x \in \Delta$ is a *typical instance* of some concept C if $x \in C^\mathcal{I}$ and there is no C-element in Δ *more typical* than x. In detail, $.^\mathcal{I}$ is the extension function that maps each concept C to $C^\mathcal{I} \subseteq \Delta$, and each role R to $R^\mathcal{I} \subseteq \Delta \times \Delta$. For concepts of $\mathcal{ALC}$, $C^\mathcal{I}$ is defined as usual. For the $\mathbf{T}$ operator, we have $(\mathbf{T}(C))^\mathcal{I} = Min_<(C^\mathcal{I})$. A model $\mathcal{M}$ can be equivalently defined by postulating the existence of a function $k_\mathcal{M} : \Delta \longmapsto \mathbb{N}$, where $k_\mathcal{M}$ assigns a finite rank to each world: the rank function $k_\mathcal{M}$ and $<$ can be defined from each other by letting $x < y$ iff $k_\mathcal{M}(x) < k_\mathcal{M}(y)$.

Given standard definitions of satisfiability of a KB in a model, we define a notion of entailment in $\mathcal{ALC} + \mathbf{T_R}$. Given a query F (either an inclusion $C \sqsubseteq D$ or an assertion $C(a)$ or an assertion of the form $R(a, b)$), we say that F is entailed from a KB in $\mathcal{ALC} + \mathbf{T_R}$, written KB $\models_{\mathcal{ALC} + \mathbf{T_R}} F$, if F holds in all $\mathcal{ALC} + \mathbf{T_R}$ models satisfying KB.

2.2   The nonmonotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$

Even if the typicality operator $\mathbf{T}$ itself is nonmonotonic (i.e. $\mathbf{T}(C) \sqsubseteq E$ does not imply $\mathbf{T}(C \sqcap D) \sqsubseteq E$), what is inferred from a KB can still be inferred from any KB' with KB $\subseteq$ KB', i.e. the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ is monotonic. In order to perform useful nonmonotonic inferences, in [14] the authors have strengthened the above semantics by restricting entailment to a class of minimal models. Intuitively, the idea is to restrict entailment to models that *minimize the untypical instances of a concept*. The resulting logic is called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ and it corresponds to a notion of *rational closure* on top of $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$. Such a notion is a natural extension of the rational closure construction provided in [16] for the propositional logic. Here below we recall the semantics of the DL $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$: details about the construction of the rational closure and the correspondence between semantics and construction can be found in [14].

The nonmonotonic semantics of $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ relies on minimal rational models that minimize the *rank of domain elements*. Informally, given two models of KB, one in which a given domain element x has rank 2 (because for instance $z < y < x$), and another in which it has rank 1 (because only $y < x$), we prefer the latter, as in this model the element x is assumed to be "more typical" than in the former. Query entailment is then restricted to minimal *canonical models*. The intuition is that a canonical model contains all the individuals that enjoy properties that are consistent with the knowledge base. This is needed when reasoning about the rank of the concepts: it is important to have them all represented. A model $\mathcal{M}$ is a minimal canonical model of KB if it satisfies KB, it is minimal and it is canonical[1]. Finally, a query F is minimally entailed from a KB (or, equivalently, F belongs to the rational closure of KB), written KB $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}}$ F, if it holds in all minimal canonical models of KB minimally satisfying $\mathcal{A}$. In [14] it is shown that query entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, i.e. the problem of checking whether a query F is in the rational closure of KB, namely that KB $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}}$ F, is in ExpTime.

3   The Logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$: Reasoning About Surprising Scenarios

In this section we define an alternative semantics that allows us to express a degree of expectedness for the typicality inclusions and to limit the number of typicality assumptions in the ABox in order to obtain less predictable scenarios. The basic idea is similar to the one proposed in [10], where a completion of an $\mathcal{ALC}+\mathbf{T}$ ABox is proposed in order to assume that every individual constant of the ABox is a typical element of the most specific concept he belongs to, if this is consistent with the knowledge base. Here we propose a similar, algorithmic construction in order to compute only *some* assumptions of typicality of domain elements/individual constants, in order to describe an alternative, surprising but plausible scenario. Constraints about the cardinality of the extensions of concepts are also introduced in order to *filter* scenarios, allowing to define *eligible* extensions of the ABox satisfying such constraints, and entailment is restricted to *minimal* scenarios, called *perfect* extensions, with respect to an order relation among

---

[1] In Theorem 10 in [14] the authors have shown that for any KB there exists a finite minimal canonical model of KB minimally satisfying the ABox.

extensions: intuitively, an extension is preferred to another one if it represents a more surprising scenario.

As mentioned above, the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$ allows to express cardinality restrictions in the TBox. More expressive DLs allow to specify (un)qualified number restrictions, in order to specify the number of possible elements filling a given role R. As an example, number restrictions allow to express that a student attends to 3 courses. Number restrictions are therefore "localized to the fillers of one particular role" [17], for instance we can have $Student \sqsubseteq = 3\, Attends.Course$ as a restriction on the number of role fillers of the role $Attends$. However one could need to express *global* restrictions on the number of domain elements belonging to a given concept, for instance to express that in the whole domain there are exactly 3 courses. In DLs not allowing cardinality restrictions one can only express that every student must attend to three courses, but not that all must attend to the same ones. In the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$, cardinality restrictions on concepts are added to the TBox as in Definition 1. They are expressions of the form either $(\geq n\ C)$ or $(\leq n\ C)$ or $(= n\ C)$, where n is a positive integer and C is a concept. This is formally defined in the next definition, where, given a set S, we write $\sharp S$ to mean the cardinality of S.

Let us first define the language $\mathcal{L}$ of this new Description Logic called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$. It is known that axioms expressing cardinality restrictions are even more expressive than inclusions $C \sqsubseteq D$ or $C \doteq D$, the last one shortening for the pair of inclusions $C \sqsubseteq D$ and $D \sqsubseteq C$ (thus expressing that the concepts C and D have the same extensions, i.e. $C^{\mathcal{I}} = D^{\mathcal{I}}$). Indeed, $C \doteq D$ means that the set of domain elements that are Cs but not Ds is empty, and viceversa (the set of domain elements that are Ds but not Cs is empty). This can be expressed by the following cardinality restriction: $(= 0\, ((C \sqcap \neg D) \sqcup (\neg C \sqcap D)))$. The same for an inclusion of the form $C \sqsubseteq D$, whose meaning is that every C element is also a D element, that can be expressed by $(= 0\, (C \sqcap \neg D))$ (the intersection of $C^{\mathcal{I}}$ and the complement of $D^{\mathcal{I}}$ is empty). Therefore, we could restrict our language to TBoxes only containing cardinality restriction, however we have decided to consider the extended language of Definition 1 for the sake of readability.

Definition 1. *We consider an alphabet of concept names $\mathcal{C}$, of role names $\mathcal{R}$, and of individual constants $\mathcal{O}$. Given $A \in \mathcal{C}$ and $R \in \mathcal{R}$, we define:*

$$C := A \mid \top \mid \bot \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C$$

An $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{exp}}$ knowledge base is a pair $(\mathcal{T}, \mathcal{A})$. $\mathcal{T}$ contains axioms of the form:

- $C \sqsubseteq C$;
- $\mathbf{T}(C) \sqsubseteq_d C$, *where $d \in \mathbb{N}^+$ is called the* degree of expectedness;
- $(\odot\ n\ C)$, *where $\odot \in \{=, \leq, \geq\}$ and $n \in \mathbb{N}^+$;*

$\mathcal{A}$ contains assertions of the form $C(a)$ and $R(a, b)$, where $a, b \in \mathcal{O}$.

Given an inclusion $\mathbf{T}(C) \sqsubseteq_d D$, the higher the degree of expectedness the more the inclusion is, in some sense, "obvious"/not surprising. Given another inclusion $\mathbf{T}(C') \sqsubseteq_{d'} D'$, with $d' < d$, we assume that this inclusion is less "obvious", more surprising with respect to the other one. As an example, let KB contain:

$$\mathbf{T}(Student) \sqsubseteq_4 SocialNetworkUser$$
$$\mathbf{T}(Student) \sqsubseteq_2 PartyParticipant$$

representing that typical students make use of social networks, and that normally they go to parties; however, the second inclusion is less obvious with respect to the first one. In other words, one can think of representing the fact that both are properties of a prototypical student, however there are more exceptions of students not taking part to parties with respect to the number of exceptions of students not being part of the social media ecosphere.

It is worth noticing that using positive integers for expressing degrees of expectedness is only a way of formalizing a partial order among typicality inclusions, however all properties expressed by typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$ are *typical* properties, even if n is high: the ontology engineer has still to distinguish between properties that are prototypical (even with some exceptions) and those that are not and do not deserve to be represented by a typicality inclusion. It is also worth noticing that degrees of expectedness are not intended to represent priorities among inclusions (as in circumscribed KBs), since specificity is provided for free by the preferential semantics of the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$.

3.1 Extensions of ABox

Given a KB, we define the finite set $\mathbb{C}$ of concepts occurring in the scope of the typicality operator, i.e. $\mathbb{C} = \{C \mid \mathbf{T}(C) \sqsubseteq_d D \in \text{KB}\}$. These are the concepts whose atypical instances we want to minimize.

Given an individual a explicitly named in the ABox, we define the set of "plausible" typicality assumptions $\mathbf{T}(C)(a)$ that can be minimally entailed from KB *without cardinality restrictions* in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, with $C \in \mathbb{C}$. We then consider an ordered set of pairs (a, C) of all possible assumptions $\mathbf{T}(C)(a)$, for all concepts $C \in \mathbb{C}$ and all individual constants a occurring in the ABox. This is formally stated in the next definition:

Definition 2 (Assumptions in $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathbf{exp}}$). *Given an $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}$ KB=$(\mathcal{T} \cup \mathcal{T}_{card}, \mathcal{A})$, where $\mathcal{T}_{card}$ is a set of cardinality restrictions and $\mathcal{T}$ does not contain cardinality restrictions, let $\mathcal{T}'$ be the set of inclusions of $\mathcal{T}$ without degrees of expectedness. Given a finite set of concepts $\mathbb{C}$, we define, for each individual name a occurring in $\mathcal{A}$:*

$$\mathbb{C}_a = \{C \in \mathbb{C} \mid (\mathcal{T}', \mathcal{A}) \models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} \mathbf{T}(C)(a)\}$$

We also define $\mathbb{C}_{\mathcal{A}} = \{(a, C) \mid C \in \mathbb{C}_a \text{ and } a \text{ occurs in } \mathcal{A}\}$ and we impose an order on the elements of $\mathbb{C}_{\mathcal{A}} = [(a_1, C_1), (a_2, C_2), \dots, (a_n, C_n)]$. Furthermore, we define the ordered multiset $d_{\mathcal{A}} = [d_1, d_2, \dots, d_n]$ *respecting the order imposed on $\mathbb{C}_{\mathcal{A}}$, where $d_i = avg(\{d \in \mathbb{N}^+ \mid \mathbf{T}(C_i) \sqsubseteq_d D \in \mathcal{T}_{card}\})$.*

Intuitively, the ordered multiset $d_{\mathcal{A}}$ is a tuple of the form $[d_1, d_2, \dots, d_n]$, where d_i is the degree of expectedness of the assumption $\mathbf{T}(C)(a)$, such that $(a, C) \in \mathbb{C}_{\mathcal{A}}$ at

position i. d_i corresponds to the average[2] of all the degrees d of typicality inclusions $\mathbf{T}(C) \sqsubseteq_d D$ in the TBox.

In order to define alternative scenarios, where not all plausible assumptions are taken into account, we consider different extensions of the ABox and we introduce an order among them, allowing to range from unpredictable to trivial ones. Starting from $d_{\mathcal{A}} = [d_1, d_2, \ldots, d_n]$, the first step is to build all alternative tuples where 0 is used in place of some d_i to represent that the corresponding typicality assertion $\mathbf{T}(C)(a)$ is no longer assumed (Definition 3). Furthermore, we define the *extension* of the ABox corresponding to a string so obtained (Definition 4).

Definition 3 (Strings of plausible assumptions $\mathbb{S}$). *Given a KB=$(\mathcal{T}, \mathcal{A})$ and the set $\mathbb{C}_{\mathcal{A}}$, let $d_{\mathcal{A}} = [d_1, d_2, \ldots, d_n]$ be as in Definition 2. We define the set $\mathbb{S}$ of all the* strings *of plausible assumptions with respect to KB as*

$$\mathbb{S} = \{[s_1, s_2, \ldots, s_n] \mid \forall i = 1, 2, \ldots, n \text{ either } s_i = d_i \text{ or } s_i = 0\}$$

Definition 4 (Extension of the ABox). *Let KB=$(\mathcal{T}, \mathcal{A})$ and let $\mathbb{C}_{\mathcal{A}} = [(a_1, C_1), (a_2, C_2), \ldots, (a_n, C_n)]$ as in Definition 2. Given a string of plausible assumptions $[s_1, s_2, \ldots, s_n] \in \mathbb{S}$ of Definition 3, we define the extension $\widehat{\mathcal{A}}$ of $\mathcal{A}$ with respect to $\mathbb{C}_{\mathcal{A}}$ and $\mathbb{S}$*

$$\widehat{\mathcal{A}} = \{\mathbf{T}(C_i)(a_i) \mid (a_i, C_i) \in \mathbb{C}_{\mathcal{A}} \text{ and } s_i \neq 0\}$$

It is easy to observe that, in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, the set of typicality assumptions that can be inferred from a KB corresponds to the extension of $\mathcal{A}$ corresponding to the string $d_{\mathcal{A}}$, that is to say no element is set to 0: all the typicality assertions of individuals occurring in the ABox, that are consistent with the KB, are assumed. On the contrary, in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$, no typicality assumption can be derived from a KB, and this corresponds to extending $\mathcal{A}$ by the assertions corresponding to the string $[0, 0, \ldots, 0]$, i.e. by the empty set.

3.2 Cardinality restrictions and perfect extensions

Let us now introduce models of the Description Logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\text{exp}}$ taking cardinality restrictions into account, as well as the notion of *eligible* extension of the ABox as a set of typicality assumptions satisfying cardinality restrictions.

Definition 5. *Given a model $\mathcal{M} = \langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$, it satisfies:*

- *(TBox)*
  - *an inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$;*
  - *a typicality inclusion $\mathbf{T}(C) \sqsubseteq_d D$ if $Min_<(C^{\mathcal{I}}) \subseteq D^{\mathcal{I}}$;*
  - *a cardinality restriction of the form $(\odot\, n\, C)$ if $\sharp C^{\mathcal{I}} \odot n$, where $\odot \in \{\leq, \geq, =\}$ and $n \in \mathbb{N}^+$;*

---

[2] Other aggregation functions could be used in order to define d_i (maximun degree, minimum degree). We aim at studying the impact of the choice on the reasoning machinery in future research.

– *(ABox)*
  - *an assertion of the form $C(a)$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$;*
  - *an assertion of the form $R(a,b)$ if $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$.*

Given a KB=$(\mathcal{T}, \mathcal{A})$, we say that a model $\mathcal{M}$ satisfies KB if it satisfies all the inclusions in $\mathcal{T}$ and all the assertions in $\mathcal{A}$.

Given a KB=$(\mathcal{T}, \mathcal{A})$, we say that an extension of $\mathcal{A}$ is an *eligible extension* if it admits a model as in Definition 5:

Definition 6 (Eligible extension $\widehat{\mathcal{A}}$). *Given an $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$ KB=$(\mathcal{T}, \mathcal{A})$ and an extension $\widehat{\mathcal{A}}$ of $\mathcal{A}$ as in Definition 4, we say that $\widehat{\mathcal{A}}$ is* eligible *if there exists an $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$ model $\mathcal{M}$ as in Definition 5 that satisfies KB'=$(\mathcal{T}, \mathcal{A} \cup \widehat{\mathcal{A}})$.*

Let us now introduce an order relation among the strings of $\mathbb{S}$ (Definition 3), corresponding to eligible extensions of the ABox:

Definition 7 (Order between eligible extensions). *Given KB=$(\mathcal{T}, \mathcal{A})$ and the set $\mathbb{S}$ of Definition 3, let $s = [s_1, s_2, \ldots, s_n]$ and $r = [r_1, r_2, \ldots, r_n]$, with $s, r \in \mathbb{S}$. Let $\widehat{\mathcal{A}}_s$ and $\widehat{\mathcal{A}}_r$ be two eligible extensions of $\mathcal{A}$ corresponding to s and r (Definition 4). We say that $s < r$ if there exists a bijection δ between s and r such that, for each $(s_i, r_j) \in \delta$, it holds that $s_i \leq r_j$, and there is at least one $(s_i, r_j) \in \delta$ such that $s_i < r_j$. We say that $\widehat{\mathcal{A}}_s$ is* more surprising *(or* less trivial*) than $\widehat{\mathcal{A}}_r$ if $s < r$.*

Intuitively, a string s whose elements are "lower" than the ones of another string r corresponds to a less trivial ABox. For instance, let us consider a KB whose typicality inclusions are $\mathbf{T}(C) \sqsubseteq_1 D$ and $\mathbf{T}(E) \sqsubseteq_2 F$, and such that $\mathbf{T}(C)(a)$, $\mathbf{T}(C)(b)$, and $\mathbf{T}(E)(b)$ are entailed in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$. Given the strings $s = [1, 1, 0]$ and $r = [1, 0, 2]$, we have that $s < r$, because there exists a bijection $\{(1, 1), (0, 0), (1, 2)\}$. The assumptions $\mathbf{T}(C)(a)$ and $\mathbf{T}(C)(b)$ corresponding to s are then considered less trivial than $\mathbf{T}(C)(a)$ and $\mathbf{T}(E)(b)$ corresponding to r. It is worth noticing that the order of Definition 7 is partial: as an example, the strings $[1, 1, 0]$ and $[0, 0, 2]$ are not comparable, in the sense that $[1, 1, 0] \not< [0, 0, 2]$ and $[0, 0, 2] \not< [1, 1, 0]$. In order to choose between two incomparable situations, we introduce the following notion of weak order: given two incomparable extensions $\widehat{\mathcal{A}}_s$ and $\widehat{\mathcal{A}}_r$, we assume that $\widehat{\mathcal{A}}_s$ is weakly less trivial than $\widehat{\mathcal{A}}_r$ if $\widehat{\mathcal{A}}_r$ is strictly included in another *eligible* extension $\widehat{\mathcal{A}}_u$ more trivial than $\widehat{\mathcal{A}}_s$, i.e. $\widehat{\mathcal{A}}_r \subset \widehat{\mathcal{A}}_u$ and $s < u$.

Definition 8 (Minimal (perfect) extensions). *Given a KB=$(\mathcal{T}, \mathcal{A})$ and the set $\mathbb{S}$ of strings of plausible assumptions (Definition 3), we say that an eligible extension $\widehat{\mathcal{A}}_s$ is* minimal *or* perfect *if there is no other eligible extension $\widehat{\mathcal{A}}_r$ which is* (weakly) more surprising *(or* (weakly) less trivial*) than $\widehat{\mathcal{A}}_s$.*

Given the above definitions, we can define a notion of entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$. Intuitively, given a query F, we check whether F follows in the monotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ from a given KB, whose ABox is augmented with extensions that are minimal (perfect) as in Definition 8. We can reason either in a skeptical way, by asking that F is

entailed if it follows in *all* KBs, obtained by considering each minimal extension of the ABox, or in a credulous way, by assuming that F is entailed if there exists at least one extension of the ABox allowing such inference. This is stated in a rigorous manner by the following definition:

Definition 9 (Entailment in $\mathcal{ALC} + \mathbf{T_R^{exp}}$). *Given a KB=$(\mathcal{T}, \mathcal{A})$ and given $\mathbb{C}$ a set of concepts, let $\mathcal{E}$ the set of all extensions of $\mathcal{A}$ that are minimal as in Definition 8. Given a query F, we say that (i) F is skeptically entailed from KB in $\mathcal{ALC} + \mathbf{T_R^{exp}}$, written $KB \models^{sk}_{\mathcal{ALC}+\mathbf{T_R^{exp}}} F$, if $(\mathcal{T}, \mathcal{A} \cup \widehat{\mathcal{A}}) \models_{\mathcal{ALC}+\mathbf{T_R}} F$ for all $\widehat{\mathcal{A}} \in \mathcal{E}$; (ii) F is credulously entailed from KB in $\mathcal{ALC} + \mathbf{T_R^{exp}}$, written $KB \models^{cr}_{\mathcal{ALC}+\mathbf{T_R^{exp}}} F$, if there exists $\widehat{\mathcal{A}} \in \mathcal{E}$ such that $(\mathcal{T}, \mathcal{A} \cup \widehat{\mathcal{A}}) \models_{\mathcal{ALC}+\mathbf{T_R}} F$.*

At a first glance, one could have the impression that the notions of rank in the semantics of $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$, where elements with lowest rank are the most typical ones, and the semantics of expectedness of Definitions 7 and 8, where lower ranks correspond to more surprising scenarios, are in conflict. However, this is not the case: ranks in the semantics are introduced in order to define the extension of typicality concepts, and this notion is also considered in the expectation semantics to select plausible typicality assumptions. The rank among extensions is rather used in order to choose surprising scenarios, to restrict the number of typicality assumptions to satisfy cardinality restrictions: the *unexpectedness* is the additional ingredient to select surprising scenarios by fixing cardinality restrictions, where all candidates try to maximize the typicality of individuals.

Let us conclude by showing an example of reasoning in the logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$.

Example 1 (Mysterious medical diagnosis). In this example we exploit the logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$ in order to provide a mysterious medical diagnosis, as an alternative to the most likely explanation, for a patient characterized by mood swings. The idea is to support the medical staff whenever the "standard" diagnosis fails, suggesting surprising alternatives that could be taken into account for further investigations.

Let KB=$(\mathcal{T}, \mathcal{A})$ as follows and let $\mathcal{T}$ be:

$$\mathbf{T}(Cancer) \sqsubseteq_1 MoodSwingsCause \tag{T1}$$
$$\mathbf{T}(BrainDisorder) \sqsubseteq_3 MoodSwingsCause \tag{T2}$$
$$\mathbf{T}(MajorDepressionAtypicalFeatures) \sqsubseteq_2 MoodSwingsCause \tag{T3}$$
$$(= 1\ MoodSwingsCause) \tag{T4}$$

the last one stating that we are interested in finding exactly one disease being responsible for mood reactivity. Concerning typicality inclusions, we represent the facts that normally, cancer causes mood swings (T1), however this admits more exceptions with respect to the fact that mood reactivity is a typical symptom of depressive disorders with atypical features (T3), which is in turn more surprising than the most "obvious" inclusion, namely that brain disorders normally cause changes in the mood of patients (T2). The cardinality restriction (T4) imposes that exactly one disease is the reason why the patient's mood swings. Let $\mathcal{A}$ be:

$Cancer(prostaticCancer)$
$BrainDisorder(bipolarDisorder)$
$MajorDepressionAtypicalFeatures(bipolarDisorder)$
$MajorDepressionAtypicalFeatures(atypicalDepression)$

In the logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$, we can infer that prostatic cancer is responsible of mood swings in our patient:

$\text{KB} \models_{\mathcal{ALC}+\mathbf{T_R^{exp}}}^{\text{sk}} MoodSwingsCause(prostaticCancer)$
$\text{KB} \models_{\mathcal{ALC}+\mathbf{T_R^{exp}}}^{\text{cr}} MoodSwingsCause(prostaticCancer)$

since there is only one perfect extension. Let

$$\mathbb{C} = \{Cancer, MajorDepressionAtypicalFeatures, BrainDisorder\}.$$

By Definition 2 above, we have that:

$\mathbb{C}_{prostaticCancer} = \{Cancer\},$
$\mathbb{C}_{atypicalDepression} = \{MajorDepressionAtypicalFeatures\},$
$\mathbb{C}_{bipolarDisorder} = \{BrainDisorder, MajorDepressionAtypicalFeatures\}$

and, obviously, $\mathbb{C}_{\mathcal{A}} = \mathbb{C}_{prostaticCancer} \cup \mathbb{C}_{atypicalDepression} \cup \mathbb{C}_{bipolarDisorder}$. Concerning the degrees of expectedness, we have $d_{\mathcal{A}} = [1, 3, 2, 2]$. As mentioned above, in $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$ the minimal model semantics forces all the consistent typicality assumptions, namely we are considering an ABox extended with the following facts:

$\mathbf{T}(Cancer)(prostaticCancer)$
$\mathbf{T}(BrainDisorder)(bipolarDisorder)$
$\mathbf{T}(MajorDepressionAtypicalFeatures)(bipolarDisorder)$
$\mathbf{T}(MajorDepressionAtypicalFeatures)(atypicalDepression)$

corresponding (in the sense of Definition 4) to the multiset $[1, 3, 2, 2]$. However, in $\mathcal{ALC} + \mathbf{T_R}$ we obtain that prostatic cancer, bipolar disorder and atypical depression are all responsible of mood reactivity in the patient, against the fact that we want to focus on only one disease: the extension corresponding to $[1, 3, 2, 2]$ is indeed not eligible in the sense of Definition 6. In order to find only one non-trivial diagnosis justifying mood swings, we consider the set $\mathbb{S}$ of all plausible strings of typicality assumptions (Definition 3). The only eligible extensions of the ABox are:

$\widehat{\mathcal{A}_1} = \{\mathbf{T}(BrainDisorder)(bipolarDisorder)\}$, by $[0, 3, 0, 0]$
$\widehat{\mathcal{A}_2} = \{\mathbf{T}(MajorDepressionAtypicalFeatures)(atypicalDepression)\}$, by $[0, 0, 0, 2]$
$\widehat{\mathcal{A}_3} = \{\mathbf{T}(Cancer)(prostaticCancer)\}$, $[1, 0, 0, 0]$
$\widehat{\mathcal{A}_4} = \{\mathbf{T}(MajorDepressionAtypicalFeatures)(bipolarDisorder)\}$, $[0, 0, 2, 0]$
$\widehat{\mathcal{A}_5} = \{\mathbf{T}(BrainDisorder)(bipolarDisorder), \mathbf{T}(MajorDepressionAtypicalFeatures)$ $(bipolarDisorder)\}$, corresponding to $[0, 3, 2, 0]$

We have that $\widehat{\mathcal{A}_1}$, $\widehat{\mathcal{A}_2}$ and $\widehat{\mathcal{A}_4}$ are less trivial than $\widehat{\mathcal{A}_5}$, because $[0, 3, 0, 0] < [0, 3, 2, 0]$, as well as $[0, 0, 0, 2] < [0, 3, 2, 0]$ and $[0, 0, 2, 0] < [0, 3, 2, 0]$. However, $\widehat{\mathcal{A}_3}$ is less trivial than $\widehat{\mathcal{A}_1}$, $\widehat{\mathcal{A}_2}$ and $\widehat{\mathcal{A}_4}$, since $[1, 0, 0, 0] < [0, 3, 0, 0]$, as well as $[1, 0, 0, 0] < [0, 0, 0, 2]$ and $[1, 0, 0, 0] < [0, 0, 2, 0]$. This allows to suggest that prostatic cancer could be a "mysterios diagnosis" for the patient having mood swings, and such a non trivial diagnosis could be confirmed by an evaluation of other typical symptoms of such a disease, e.g. nocturia.

4 A Decision Procedure for Reasoning in $\mathcal{ALC} + \mathbf{T_R^{exp}}$

In this section we describe a decision procedure for reasoning in the logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$. We consider skeptical and credulous entailment. In both cases, we exploit the decision procedure to show that the problem of entailment in the logic $\mathcal{ALC} + \mathbf{T_R^{exp}}$ is in ExpTime. This allows us to conclude that reasoning about typicality and defeasible inheritance in surprising scenarios is essentially inexpensive, in the sense that reasoning retains the same complexity class of the underlying standard Description Logic $\mathcal{ALC}$, which is known to be ExpTime complete [19].

Following the definitions of nonmonotonic entailment introduced in the previous section, given an $\mathcal{ALC} + \mathbf{T_R^{exp}}$ KB=$(\mathcal{T}, \mathcal{A})$ and a query F, we define a procedure computing the following three steps:

1. compute the set $\mathbb{C}_a$ of all typicality assumptions that are minimally entailed from the knowledge base in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$;
2. compute all possible extensions of the ABox and select perfect extensions;
3. check whether the query F is entailed from at least one extension/all the extensions of KB in the monotonic logic $\mathcal{ALC} + \mathbf{T_R}$ plus cardinality restrictions.

Step 3 is based on reasoning in the monotonic logic $\mathcal{ALC} + \mathbf{T_R}$: to this aim, the procedure relies on a polynomial encoding of $\mathcal{ALC} + \mathbf{T_R}$ into $\mathcal{ALC}$ introduced in [20] and then on reasoning with cardinality restrictions. Step 1 is based on reasoning in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$: in this case, the procedure computes the rational closure of an $\mathcal{ALC} + \mathbf{T_R}$ knowledge base by means of the algorithm introduced in [14], which is sound and complete with respect to the minimal model semantics recalled in Section 2.2. Also the algorithm computing the rational closure relies on reasoning in the monotonic logic $\mathcal{ALC} + \mathbf{T_R}$, then on the above mentioned polynomial encoding in $\mathcal{ALC}$. We assume unary encoding of numbers in cardinality restrictions in order to exploit the results in [21], namely that reasoning in $\mathcal{ALCO}$, extending $\mathcal{ALC}$ with qualified number restrictions, is ExpTime-complete also with cardinality restrictions. Due to space limitations, here we only introduce the overall procedure for reasoning in $\mathcal{ALC} + \mathbf{T_R^{exp}}$ and we analyze its complexity, whereas we remind to the accompanying report [22] for the procedures for reasoning in $\mathcal{ALC} + \mathbf{T_R}$ and $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$.

Let KB=$(\mathcal{T} \cup \mathcal{T}_{card}, \mathcal{A})$ be an $\mathcal{ALC} + \mathbf{T_R^{exp}}$ knowledge base, where $\mathcal{T}_{card}$ is a set of cardinality restrictions and $\mathcal{T}$ does not contain cardinality restrictions. Let $\mathcal{T}'$ be the set of inclusions of $\mathcal{T}$ without the degrees of expectedness: $\mathcal{T}' = \{C \sqsubseteq D \mid C \sqsubseteq_n D \in \mathcal{T}\}$, that the procedure will take into account in order to reason in $\mathcal{ALC} + \mathbf{T_R}$ and $\mathcal{ALC} + \mathbf{T_R^{RaCl}}$ for checking query entailment and finding all plausible typicality

assumptions, respectively. Other inputs of the procedure are a finite set of concepts $\mathbb{C}$ and a query F. Algorithm 1 checks whether F is skeptically entailed from the KB in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$, namely whether $\text{KB} \models^{sk}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$.

Algorithm 1 Skeptical entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$: $\text{KB} \models^{sk}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$

1: **procedure** SKEPTICALENTAILMENT$((\mathcal{T} \cup \mathcal{T}_{card}, \mathcal{A}), \mathcal{T}', F, \mathbb{C})$
2: $\quad \mathbb{C}_{\mathcal{A}} \leftarrow \emptyset$ $\qquad\qquad\qquad\qquad\qquad\qquad$ ▷ *build the set $\mathbb{S}$ of plausible assumptions*
3: $\quad$ **for** *each $C \in \mathbb{C}$* **do**
4: $\qquad$ **for** *each individual $a \in \mathcal{A}$* **do** $\qquad\qquad\qquad$ ▷ *Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$*
5: $\qquad\quad$ **if** $(\mathcal{T}', \mathcal{A}) \models_{\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}} \mathbf{T}(C)(a)$ **then** $\mathbb{C}_{\mathcal{A}} \leftarrow \mathbb{C}_{\mathcal{A}} \cup \{\mathbf{T}(C)(a)\}$
6: $\quad d_{\mathcal{A}} \leftarrow$ build the ordered multiset of avg degrees of Definition 2 given $\mathcal{T}$ and $\mathbb{C}_{\mathcal{A}}$
7: $\quad \mathbb{S} \leftarrow$ build strings of plausible extensions as in Definition 3 given $\mathbb{C}_{\mathcal{A}}$ and $d_{\mathcal{A}}$
8: $\quad \mathcal{A}_{pl} \leftarrow \emptyset$ $\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ ▷ *build plausible extensions of $\mathcal{A}$*
9: $\quad$ **for** *each $d_i \in \mathbb{S}$* **do**
10: $\qquad$ build the extension $\widehat{\mathcal{A}}_i$ corresponding to d_i
11: $\qquad \mathcal{A}_{pl} \leftarrow \mathcal{A}_{pl} \cup \widehat{\mathcal{A}}_i$
12: $\quad \mathcal{A}_{el} \leftarrow \emptyset$ $\qquad\qquad$ ▷ *select eligible extensions checking cardinality restrictions*
13: $\quad$ **for** *each $\widehat{\mathcal{A}}_i \in \mathcal{A}_{pl}$* **do** $\qquad$ ▷ *Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ plus cardinality restrictions*
14: $\qquad$ **if** $(\mathcal{T}' \cup \mathcal{T}_{card}, \mathcal{A} \cup \widehat{\mathcal{A}}_i)$ *is satisfiable in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$* **then**
15: $\qquad\quad \mathcal{A}_{el} \leftarrow \mathcal{A}_{el} \cup \widehat{\mathcal{A}}_i$
16: $\quad$ **for** *each $\widehat{\mathcal{A}}_i \in \mathcal{A}_{el}$* **do** $\qquad\qquad\qquad$ ▷ *check preference among extensions of $\mathcal{A}$*
17: $\qquad$ **for** *each $\widehat{\mathcal{A}}_j \in \mathcal{A}_{el}$* **do**
18: $\qquad\quad$ **if** $d_i \leq d_j$ **then** *let* $\widehat{\mathcal{A}}_i < \widehat{\mathcal{A}}_j$
19: $\quad \mathcal{E} \leftarrow \{\widehat{\mathcal{A}}_i \mid \nexists \widehat{\mathcal{A}}_j \in \mathcal{A}_{el} \text{ such that } \widehat{\mathcal{A}}_j < \widehat{\mathcal{A}}_i\}$ $\qquad$ ▷ *select perfect extensions*
20: $\quad$ **for** *each $\widehat{\mathcal{A}}_i \in \mathcal{E}$* **do** $\quad$ ▷ *query entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ plus cardinality restrictions*
21: $\qquad$ **if** $(\mathcal{T}' \cup \mathcal{T}_{card}, \mathcal{A} \cup \widehat{\mathcal{A}}_i) \not\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}} F$ **then**
22: $\qquad\quad$ **return** $KB \not\models^{sk}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$ $\qquad\qquad$ ▷ *a perfect extension not entailing F*
23: $\quad$ **return** $KB \models^{sk}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$ $\qquad\qquad$ ▷ *F is entailed in all perfect extensions*

In order to check whether F is credulously entailed from the KB, that is to say $\text{KB} \models^{cr}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$, the algorithm is obtained by replacing lines 20-23 in Algorithm 1 by the following ones:

20: $\quad$ **for** each $\widehat{\mathcal{A}}_i \in \mathcal{E}$ **do**
21: $\qquad$ **if** $(\mathcal{T}' \cup \mathcal{T}_{card}, \mathcal{A} \cup \widehat{\mathcal{A}}_i) \models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}} F$ **then**
22: $\qquad\quad$ **return** $\text{KB} \models^{cr}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$
23: $\quad$ **return** $\text{KB} \not\models^{cr}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{exp}} F$

By exploiting the procedures above, we show that:

Theorem 1 (Complexity of entailment). *Given a KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{exp}$ and a query F whose size is polynomial in the size of KB, assuming the unary encoding of numbers in*

cardinality restrictions of KB, the problem of checking skeptically (resp. credulously) whether KB $\models^{sk}_{\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}} F$ (resp. KB $\models^{cr}_{\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}} F$) is EXPTIME-*complete.*

Proof. (*sketch, see* [22] *for the complete proof*) The algorithm checks, for each concept $C \in \mathbb{C}$ and for each individual name a whether $\mathbf{T}(C)(a)$ is minimally entailed from KB in the nonmonotonic logic $\mathcal{ALC}+\mathbf{T}^{RaCl}_{\mathbf{R}}$. By definition, the size of $\mathbb{C}$ is $O(n)$. For each $\mathbf{T}(C)(a)$ (they are $O(n^2)$) the algorithm relies on reasoning in $\mathcal{ALC}+\mathbf{T}^{RaCl}_{\mathbf{R}}$, which is in EXPTIME. Building $d_{\mathcal{A}}$ can be solved with $O(n^2)$ operations. For building the set $\mathbb{S}$ of plausible extensions we have to consider all possible strings obtained by assuming (or not) each typicality assumption $\mathbf{T}(C)(a)$, that are $O(n^2)$: for each d_i, we have two options ($d_i = 0$ or $d_i \neq 0$) , then $2 \times 2 \times \ldots \times 2$ different strings, thus $\mathbb{S}$ has exponential size in n. Checking, for each extension of $\mathcal{A}$, if it satisfies cardinality constraints requires $O(2^{n^2})$ calls to satisfiability in $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}$ plus cardinality restrictions, that is in EXPTIME. Ordering extensions of $\mathcal{A}$ and finding perfect extensions can be solved in EXPTIME, then the algorithm relies on reasoning in monotonic $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}$ plus cardinality restrictions in order to check whether the query F is entailed in perfect extensions in $\mathcal{E}$, whose size is $O(2^n)$: we have $O(2^n)$ call to query entailment in $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}$, which is an EXPTIME-complete problem. □

Since reasoning in the underlying standard $\mathcal{ALC}$ is EXPTIME-complete, we can conclude that reasoning about typicality in surprising scenarios is essentially inexpensive.

5 Conclusions

In this work we have provided a nonmonotonic procedure for preferential Description Logics in order to reason about surprising scenarios. We have introduced the Description Logic of typicality $\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}$, an extension of $\mathcal{ALC}$ with a typicality operator $\mathbf{T}$ allowing to (i) express typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$, where d is a positive integer representing a degree of expectedness; (ii) reason in presence of restrictions on the cardinality of concepts. We have also described a procedure for reasoning in $\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}$ exploiting reasoning mechanisms in the logics $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}$ and $\mathcal{ALC}+\mathbf{T}^{RaCl}_{\mathbf{R}}$, the last one relying on a notion of rational closure for Description Logics. This procedure allowed us to show that entailment in $\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}$ is EXPTIME complete as the underlying $\mathcal{ALC}$, therefore it is essentially inexpensive, once unary encoding of numbers in cardinality restrictions is assumed.

The extension of DLs of typicality with cardinality restrictions is of its own interest, and one can think of considering cardinality restrictions not limited to surprising scenarios of the logic $\mathcal{ALC}+\mathbf{T}^{exp}_{\mathbf{R}}$, but directly applied to the nonmonotonic semantics of $\mathcal{ALC}+\mathbf{T}^{RaCl}_{\mathbf{R}}$. Furthermore, we aim at studying also cardinality restrictions on roles.

In future work we aim at extending this approach to more expressive Description Logics, in particular the logics underlying the standard language for ontology engineering OWL. As a first step, in [23] the logic with the typicality operator and the rational closure construction have been applied to the description logic $\mathcal{SHIQ}$.

A comparison with probabilistic approaches will be also object of further investigations. To the best of our knowledge, the literature lacks a formalization of surprising

scenarios in probabilistic formalizations of knowledge, however it is worth observing that a surprising scenario could be defined as a set of facts with a low probability, then one can think of restricting the attention to less probable outcomes.

Acknowledgements

References

1. Pozzato, G.L.: Preferential description logics meet sports entertainment: cardinality restrictions and perfect extensions for a better royal rumble match. In Ancona, D., Maratea, M., Mascardi, V., eds.: Proceedings of the 30th Italian Conference on Computational Logic, Genova, Italy, July 1-3, 2015. Volume 1459 of CEUR Workshop Proceedings., CEUR-WS.org (2015) 159–174
2. Bonatti, P.A., Lutz, C., Wolter, F.: The complexity of circumscription in dls. Journal of Artificial Intelligence Research (JAIR) **35** (2009) 717–773
3. Baader, F., Hollunder, B.: Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. Journal of Automated Reasoning (JAR) **15**(1) (1995) 41–68
4. Bonatti, P.A., Faella, M., Sauro, L.: Defeasible inclusions in low-complexity dls. Journal of Artificial Intelligence Research (JAIR) **42** (2011) 719–764
5. Donini, F.M., Nardi, D., Rosati, R.: Description logics of minimal knowledge and negation as failure. ACM Transactions on Computational Logics (ToCL) **3**(2) (2002) 177–225
6. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In Janhunen, T., Niemelä, I., eds.: Logics in Artificial Intelligence - Proceedings of the12th European Conference (JELIA 2010). Volume 6341 of Lecture Notes in Computer Science (LNCS)., Springer (2010) 77–90
7. Casini, G., Straccia, U.: Defeasible Inheritance-Based Description Logics. Journal of Artificial Intelligence Research (JAIR) **48** (2013) 415–473
8. Straccia, U.: Default inheritance reasoning in hybrid kl-one-style logics. In: Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI'93), Morgan Kaufmann (1993) 676–681
9. Bonatti, P.A., Faella, M., Petrova, I., Sauro, L.: A new semantics for overriding in description logics. Artificial Intelligence **222** (2015) 1–48
10. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: ALC+T: a preferential extension of description logics. Fundamenta Informaticae **96** (2009) 341–372
11. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: A NonMonotonic Description Logic for Reasoning About Typicality. Artificial Intelligence **195** (2013) 165 – 202
12. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Preferential vs Rational Description Logics: which one for Reasoning About Typicality? . In Coelho, H., Studer, R., Wooldridge, M., eds.: Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010). Volume 215 of FAIA (Frontiers in Artificial Intelligence and Applications)., Lisbon, Portugal, IOS Press (August 2010) 1069 – 1070

13. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Reasoning about typicality in low complexity DLs: the logics $\mathcal{EL}^{\perp}\mathbf{T}_{min}$ and $DL\text{-}Lite_c\mathbf{T}_{min}$. In Walsh, T., ed.: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), Barcelona, Spain, IOS Press (2011) 894–899
14. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Semantic characterization of Rational Closure: from Propositional Logic to Description Logics. Artificial Intelligence **226** (2015) 1–33
15. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence **44**(1-2) (1990) 167–207
16. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**(1) (1992) 1–60
17. Baader, F., Buchheit, M., Hollunder, B.: Cardinality restrictions on concepts. Artificial Intelligence **88**(1-2) (1996) 195–213
18. Bordino, I., Mejova, Y., Lalmas, M.: Penguins in sweaters, or serendipitous entity search on user-generated content. In He, Q., Iyengar, A., Nejdl, W., Pei, J., Rastogi, R., eds.: 22nd ACM International Conference on Information and Knowledge Management, CIKM'13, San Francisco, CA, USA, October 27 - November 1, 2013. (2013) 109–118
19. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: The Description Logic Handbook - Theory, Implementation, and Applications, 2nd edition. Cambridge (2007)
20. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Minimal model semantics and rational closure in description logics. In Eiter, T., Glimm, B., Kazakov, Y., Krötzsch, M., eds.: Informal Proceedings of the 26th International Workshop on Description Logics. Volume 1014 of CEUR Workshop Proceedings., CEUR-WS.org (2013) 168–180
21. Tobies, S.: The complexity of reasoning with cardinality restrictions and nominals in expressive description logics. Journal of Artificial Intelligence Research (JAIR) **12** (2000) 199–217
22. Pozzato, G.L.: On Reasoning about surprising scenarios in Preferential Description Logics. In: Technical Report 02/2015, http://www.di.unito.it/~pozzato/papers/RT022015.pdf, Dip. di Informatica, Univ. di Torino. (2015)
23. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Rational closure in $\mathcal{SHIQ}$. In: DL 2014, 27th International Workshop on Description Logics. Volume 1193 of CEUR Workshop Proceedings., CEUR-WS.org (2014) 543–555