On Probabilities of Exceptions in Description Logics of Typicality

Gian Luca Pozzato

Dipartimento di Informatica, Università di Torino, Italy
gianluca.pozzato@unito.it

Abstract. We describe a nonmonotonic procedure for preferential Description Logics in order to reason about typicality by taking probabilities of exceptions into account. We consider an extension, called $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$, of the logic of typicality $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ by inclusions of the form $\mathbf{T}(C) \sqsubseteq_p D$, whose intuitive meaning is that "typical Cs are Ds with a probability p". We consider a notion of extension of an ABox containing only some typicality assertions, then we equip each extension with a probability. We then restrict entailment of a query F to those extensions whose probabilities belong to a given and fixed range. We propose a decision procedure for reasoning in $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ and we exploit it to show that entailment is EXPTIME-complete as for the underlying $\mathcal{ALC}$.

1 Introduction

Description Logics [1], for short: DLs, represent one of the most important formalisms of knowledge representation and are at the base of the languages for building ontologies in the Semantic Web such as OWL. Standard Description Logics are not able to represent prototypical properties and to reason about defeasible inheritance. Recalling a well known example coming from the literature of nonmonotonic reasoning, we can have a TBox representing that birds fly ($\mathit{Bird} \sqsubseteq \mathit{Fly}$), but that penguins are birds that do not fly ($\mathit{Penguin} \sqsubseteq \mathit{Bird}$ and $\mathit{Penguin} \sqsubseteq \neg \mathit{Fly}$). This knowledge base is consistent only if there are no penguins. In order to tackle this problem, nonmonotonic extensions of Description Logics have been actively investigated since the early 90s [2–8], allowing one to represent *prototypical* properties of classes and to reason about *defasible* inheritance.

A simple but powerful nonmonotonic extension of DLs is proposed in [9]: in this approach "typical" or "normal" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying DL, and a TBox can contain inclusions of the form $\mathbf{T}(C) \sqsubseteq D$ to represent that "typical Cs are also Ds" or "normally, Cs have the property D". The Description Logic so obtained is called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ and, 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 a typical activist of the Five Stars Movement is against fraudsters, however he changes his mind if he also supports the so called "Government of changes" with The League: as an example, a typical activist supporting the government is no longer angry about the fact that the Italy's supreme court has ruled that The League has recouped some 49 million of euros of public money received by the party's former leader. This can be formalized as follows:

$$\mathbf{T}(\mathit{FiveStarsActivist}) \sqsubseteq \exists \mathit{against}.\mathit{Fraudsters}$$
$$\mathbf{T}(\mathit{FiveStarsActivist} \sqcap \mathit{GovernmentOfChangeSupporter}) \sqsubseteq \neg \exists \mathit{against}.\mathit{Fraudsters}$$

The semantics of $\mathbf{T}$ is characterized by the properties of *rational logic* [10], recognized as the core properties of nonmonotonic reasoning. As a consequence, $\mathbf{T}$ inherits well-established properties like *specificity*: in the example, if one knows that Beppenito is a typical Five Stars activist supporting the Government of Change, then the logic $\mathcal{ALC} + \mathbf{T_R}$ allows us to infer that, normally, he is not against fraudsters, giving preference to the most specific information.

The logic $\mathcal{ALC} + \mathbf{T_R}$ itself is 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 Gianpiercarloberto is a Five Stars activist, i.e. $\mathit{FiveStarsActivist}(\mathit{gianpiercarloberto})$ belongs to KB, it is not possible to infer that he is against fraudsters. This would be possible only if the stronger information that Gianpiercarloberto is a *typical* activist $\mathbf{T}(\mathit{FiveStarsActivist})(\mathit{gianpiercarloberto})$ belongs to (or can be inferred from) KB. In order to overwhelm this limit and perform useful inferences, in [11] 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 [10] 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, that we call $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, supports typicality assumptions, so that if one knows that Gianpiercarloberto is a Five Stars activist, one can nonmonotonically assume that he is also a *typical* Five Stars activist if this is consistent, and therefore that he is against fraudsters. From a semantic point of view, the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ 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 $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ imposes to consider *all* typicality assumptions that are consistent with a given KB. Following the above example about birds and penguins, if the TBox further contains that, normally, birds have wings and have a small size, if it is consistent to assume that Tweety and Kirby are typical birds, then the logic imposes to infer that they both fly, have wings and have a small size. This seems to be too strong in several application domains: it could be useful to reason about scenarios with *exceptional individuals*, or one could need to assign different *probabilities* to typicality inclusions. In the example, one could need to represent that the properties of flying, having wings and having a small size are all typical properties of birds: however, it could be needed to also describe that the probability of finding exceptional birds not having wings is lower than the one of finding exceptional birds not having a small size.

In this work we describe a new an extension of $\mathcal{ALC}$, called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, by means of typicality inclusions equipped by *probabilities of exceptionality* of the form $\mathbf{T}(C) \sqsubseteq_p D$, where $p \in (0, 1)$. The intuitive meaning is that "typical Cs are also Ds with a probability p" or "normally, Cs are Ds and the probability of having exceptional Cs – not being Ds – is $1 - p$". For instance, we can have

$\mathbf{T}(ItalianTeenAger) \sqsubseteq_{0.9} AppUser$
$\mathbf{T}(ItalianTeenAger) \sqsubseteq_{0.6} \exists listenTo.Trap$

whose intuitive meaning is that being users of apps for mobile devices and listening to trap music are both typical properties of Italian teen agers, however the probability of having exceptional teen agers not using apps is lower than the one of finding exceptional teens not listening to trap music, in particular we have the evidence that the probability of not having exceptions is 90% and 60%, respectively.

As a difference with DLs under the distributed semantics introduced in [12, 13], where probabilistic axioms of the form $p :: C \sqsubseteq D$ are used to capture uncertainty in order to represent that Cs are Ds with probability p, in the logic $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ we are able to ascribe typical properties to concepts and to reason about probabilities of exceptions to those typicalities. We define different extensions of an ABox containing only some of the "plausible" typicality assertions: each extension represents a scenario having a specific probability. Then, we provide a notion of nonmonotonic entailment restricted to extensions whose probabilities belong to a given and fixed range, in order to reason about scenarios that are not necessarily the most probable. We introduce a decision procedure for checking entailment in $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ and we exploit it in order to show that reasoning in $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ with probabilities of exceptions is EXPTIME complete, therefore we retain the same complexity of the underlying standard $\mathcal{ALC}$. This work extends and revises the preliminary results presented in [14].

2 Preferential Description Logics

The logic $\mathcal{ALC} + \mathbf{T_R}$ is obtained by adding to standard $\mathcal{ALC}$ the typicality operator $\mathbf{T}$ [9]. 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*:

Definition 1. *A* rational model $\mathcal{M}$ *is any structure* $\langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$ *where:* • $\Delta^{\mathcal{I}}$ *is a non-empty set of items called the domain;* • $<$ *is an irreflexive, transitive, well-founded and modular (for all* x, y, z *in* $\Delta^{\mathcal{I}}$*, if* $x < y$ *then either* $x < z$ *or* $z < y$*) relation over* $\Delta^{\mathcal{I}}$*;* • $.^{\mathcal{I}}$ *is the extension function that maps each concept* C *to* $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$*, and each role* R *to* $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$*. For concepts of* $\mathcal{ALC}$*,* $C^{\mathcal{I}}$ *is defined as usual. For* $\mathbf{T}$*, we have* $(\mathbf{T}(C))^{\mathcal{I}} = Min_<(C^{\mathcal{I}})$.

The intuitive idea is as follows: the preference relation among domain elements defines their level of exceptionality, namely $x < y$ means that x is "more normal" than y. Typical members of a concept C are the minimal elements of the extension $C^{\mathcal{I}}$ of C with respect to the preference relation $<$. An element $x \in \Delta^{\mathcal{I}}$ is a *typical instance* of concept C if $x \in C^{\mathcal{I}}$ and there is no C-element in $\Delta^{\mathcal{I}}$ *more typical* than x, i.e. there is no $z \in C^{\mathcal{I}}$ such that $z < x$.

A model $\mathcal{M}$ can be equivalently defined by postulating the existence of a function $k_{\mathcal{M}} : \Delta^{\mathcal{I}} \longmapsto \mathbb{N}$, where $k_{\mathcal{M}}$ assigns a finite rank to each domain element:

Definition 2 (Rank of a domain element). *Given a model* $\mathcal{M} = \langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$*, the rank* $k_{\mathcal{M}}$ *of a domain element* $x \in \Delta^{\mathcal{I}}$*, is the length of the longest chain* $x_0 < \ldots < x$ *from* x *to a minimal* x_0 *(i.e. such that there is no* x' *such that* $x' < x_0$*).*

$k_{\mathcal{M}}$ and $<$ can be defined from each other by letting $x < y$ if and only if $k_{\mathcal{M}}(x) < k_{\mathcal{M}}(y)$.

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, written KB $\models_{\mathcal{ALC}+\mathbf{T_R}} F$, if F holds in all $\mathcal{ALC} + \mathbf{T_R}$ models satisfying KB.

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 $\mathcal{ALC} + \mathbf{T_R}$ is monotonic. In order to perform useful nonmonotonic inferences, in [11] 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 $\mathcal{ALC} + \mathbf{T_R}$. Such a notion is a natural extension of the rational closure construction provided in [10] for the propositional logic.

The nonmonotonic semantics of $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ relies on minimal rational models that minimize the rank of domain elements of Definition 2 above. Intuitively, 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. Formal definitions follow.

Definition 3 (Minimal models). *Given* $\mathcal{M} = \langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$ *and* $\mathcal{M}' = \langle \Delta^{\mathcal{I}'}, <', .^{\mathcal{I}'} \rangle$ *we say that* $\mathcal{M}$ *is preferred to* $\mathcal{M}'$*, written* $\mathcal{M} < \mathcal{M}'$*, if the following conditions hold:*

- $\Delta^{\mathcal{I}} = \Delta^{\mathcal{I}'}$
- $C^{\mathcal{I}} = C^{\mathcal{I}'}$ *for all concepts* C
- *for all* $x \in \Delta^{\mathcal{I}}$*, it holds that* $k_{\mathcal{M}}(x) \leq k_{\mathcal{M}'}(x)$ *whereas there exists* $y \in \Delta^{\mathcal{I}}$ *such that* $k_{\mathcal{M}}(y) < k_{\mathcal{M}'}(y)$*.*

Given a KB, we say that $\mathcal{M}$ *is a minimal model of* K *with respect to* $<$ *if it is a model satisfying* K *and there is no* $\mathcal{M}'$ *model satisfying* K *such that* $\mathcal{M}' < \mathcal{M}$*.*

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 KB. A model $\mathcal{M}$ is a minimal canonical model of KB if it satisfies KB, it is minimal and it is canonical. In [11], Theorem 10 shows that for any consistent KB there exists a finite minimal canonical model of KB. A query F is minimally entailed from a KB, written KB $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} F$, if it holds in all minimal canonical models of KB. In [11] it is shown that query entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ is in EXPTIME.

3 Dealing with Probabilities of Exceptions: the Logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$

We introduce a semantics that allows us to equip a typicality inclusion with the probability of *not* having exceptions for that, and then to reason about such inclusions. In the resulting Description Logic, called $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, a typicality inclusion has the form

$$\mathbf{T}(C) \sqsubseteq_p D,$$

and its intuitive meaning is "normally, Cs are also Ds with probability p" or, in other words, "typical Cs are also Ds, and the probability of having exceptional Cs not being Ds is $1 - p$". We then define a nonmonotonic procedure whose aim is to describe alternative completions of the ABox obtained by assuming typicality assertions about the individuals explicitly named in the ABox: the basic idea is similar to the one proposed in [9], 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. An analogous approach is proposed in [15], where different extensions of the ABox are introduced in order to define plausible but *surprising* scenarios. Here we propose a similar, algorithmic construction in order to compute only *some* assumptions of typicality of individual constants, in order to describe alternative scenarios having different probabilities: different extensions/scenarios are obtained by considering different sets of typicality assumptions of the form $\mathbf{T}(C)(a)$, where a occurs in the ABox.

Definition 4. *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 $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ knowledge base is a pair $(\mathcal{T}, \mathcal{A})$. $\mathcal{T}$ contains axioms of the form either (i) $C \sqsubseteq C$ or (ii) $\mathbf{T}(C) \sqsubseteq_p C$, where $p \in \mathbb{R}, p \in (0, 1)$. $\mathcal{A}$ contains assertions of the form either (i) $C(a)$ or (ii) $R(a, b)$, where $a, b \in \mathcal{O}$.

Given an inclusion $\mathbf{T}(C) \sqsubseteq_p D$, the higher the probability p the more the inclusion is "exceptions-free" or, equivalently, the less is the probability of having exceptional Cs not being also Ds. In this respect, the probability p is a real number included in the open interval $(0, 1)$: the probability 1 is not allowed, in the sense that an inclusion $\mathbf{T}(C) \sqsubseteq_1 D$ (the probability of having exceptional Cs not being Ds is 0) corresponds to a *strict* inclusion $C \sqsubseteq D$ (all Cs are Ds). Given another inclusion $\mathbf{T}(C') \sqsubseteq_{p'} D'$, with $p' < p$, we assume that this inclusion is less "strict" than the other one, i.e. the probability of having exceptional C's is higher than the one of having exceptional Cs with respect to properties D' and D, respectively. Recalling the example of the Introduction, where KB contains $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_{0.9} \mathit{AppUser}$ and $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_{0.6} \exists \mathit{listenTo}.\mathit{TrapMusic}$, we have that typical Italian teen agers make use of apps for mobile devices, and that normally they also listen to trap music; however, the second inclusion is less probable with respect to the first one: both are properties of a prototypical Italian teen ager, however there is a higher probability of finding exceptions of teens not listening to trap music with respect to the probability of having Italian teen agers not using apps for mobile devices.

Before introducing formal definitions, we provide an example inspired to Example 1 in [15] in order to give an intuitive idea of what we mean for reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ with probabilities of exceptions. We will complete it with part 2 in Example 3.

Example 1 (Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ part 1). We aim at providing a formalization of some information about illnesses and symptoms. Let KB = $(\mathcal{T}, \mathcal{A})$ where $\mathcal{T}$ is:

$$AtypicalDepressed \sqsubseteq Depressed \quad (1)$$
$$\mathbf{T}(Depressed) \sqsubseteq_{0.85} \neg \exists hasSymptom.MoodReactivity \quad (2)$$
$$\mathbf{T}(AtypicalDepressed) \sqsubseteq_{0.7} \exists hasSymptom.MoodReactivity \quad (3)$$
$$\mathbf{T}(ProstateCancerPatient) \sqsubseteq_{0.6} \exists hasSymptom.MoodReactivity \quad (4)$$
$$\mathbf{T}(ProstateCancerPatient) \sqsubseteq_{0.8} \exists hasSymptom.Nocturia \quad (5)$$

The above TBox $\mathcal{T}$ represents that (2), normally, *depressed* people do not have *mood reactivity*, namely the ability to feel better temporarily in response to positive life events. On the contrary, (3) states that this is a typical symptom of a depression with atypical features, known as *atypical depression*, that shares many of the typical symptoms of depression but is characterized by improved mood in response to positive events. Inclusion (1) intuitively represents that atypical depression is a kind of depression. Mood reactivity, as well as *nocturia*, are also typical symptoms of *prostatic cancer* (inclusions (4) and (5), respectively): more in detail, (4) says that we have a probability of 60% of not having exceptional prostatic cancer patients with no mood swings, whereas (5) says that the probability of not having exceptional prostatic cancer patients without nocturia is 80% or, alternatively, probabilities of having exceptional prostatic cancer patients with no mood swing and without nocturia are of the 40% and of the 20%, respectively.

Concerning TBox reasoning, as a first example, we have that in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ we can infer[1] that, normally, depression in patients is not classified as *Atypical depression*:

$$\mathbf{T}(Depressed) \sqsubseteq \neg AtypicalDepressed,$$

and this is a wanted inference, since, in normal circumstances, members of a class do not belong to a subclass containing exceptional individuals.

As another example, we have that

$$(6)\ \mathbf{T}(Depressed \sqcap Spleenless) \sqsubseteq \neg \exists hasSymptom.MoodReactivity$$

follows from KB, and this is also a wanted inference, since undergoing spleen removal is irrelevant with respect to mood reactivity as far as we know. This is a nonmonotonic inference that does no longer follow if it is discovered that typical depressed people without their spleen are subject to mood reactivity: given

$$\mathcal{T}' = \mathcal{T} \cup \{\mathbf{T}(Depressed \sqcap Spleenless) \sqsubseteq \exists hasSymptom.MoodReactivity\},$$

we have that (6) does no longer follow from KB with $\mathcal{T}'$ in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$.

[1] At this point of the presentation we only want to give an intuition of inferences characterizing $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$. Technical details and definitions will be provided in Definition 8.

As for rational closure, the set of inclusions that are entailed from a $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ KB is closed under the property known as *rational monotonicity*: for instance, from KB and the fact that the inclusion representing that, normally, depressed people are not elder ($\mathbf{T}(\mathit{Depressed}) \sqsubseteq \neg \mathit{Elder}$) is not entailed from KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, it follows that we can infer the inclusion $\mathbf{T}(\mathit{Depressed} \sqcap \mathit{Elder}) \sqsubseteq \neg \exists \mathit{hasSymptom}.\mathit{MoodReactivity}$, namely, a typical depressed and elder patient has not mood reactivity (the subconcept $\mathit{Depressed} \sqcap \mathit{Elder}$ inherits the typical properties of the concept $\mathit{Depressed}$).

Concerning ABox reasoning, if we know that Jim is depressed, that is to say $\mathcal{A} = \{\mathit{Depressed}(\mathit{jim})\}$, then we can infer that Jim has not mood swings with a probability of 85%, since $\mathbf{T}(\mathit{Depressed})(\mathit{jim})$ is minimally entailed from KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathit{RaCl}}$ and the inclusion (2) is equipped by a probability of 0.85. If we discover that Jim is an atypical depressed, then $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ allows us to retract such inference, whereas the fact that Jim has mood swings ($\exists \mathit{hasSymptom}.\mathit{MoodReactivity}(\mathit{jim})$) is entailed and evaluated having probability of 70%. The same conclusions are also entailed in case we discover that Jim is elder, i.e. $\mathit{Elder}(\mathit{jim})$ is added to the ABox, in detail:

- from $(\mathcal{T}, \{\mathit{Depressed}(\mathit{jim}), \mathit{Elder}(\mathit{jim})\})$, the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ allows us to infer $\neg \exists \mathit{hasSymptom}.\mathit{MoodReactivity}(\mathit{jim})$ with probability of 85%;
- from $(\mathcal{T}, \{\mathit{AtypicalDepressed}(\mathit{jim}), \mathit{Elder}(\mathit{jim})\})$, the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ allows us to infer $\exists \mathit{hasSymptom}.\mathit{MoodReactivity}(\mathit{jim})$ with probability of 70%.

It is worth noticing that it is possible to have knowledge bases containing inclusions of the form $\mathbf{T}(C) \sqsubseteq_p D$, where $p \leq 0.5$ that, if wrongly interpreted, could be considered as counter intuitive. For instance, the inclusion $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_{0.3} \mathit{SportLover}$ could be wrongly interpreted as "normally, Italian teen agers do not love sport". However, probabilities in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ are not intended to express degrees of belief of the inclusions they equip. In the example, even if its corresponding probability of exceptionality is low, the right interpretation of $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_{0.3} \mathit{SportLover}$ is that loving sport is *anyway* a property of a prototypical Italian teen ager: as a difference with $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_{0.9} \mathit{AppUser}$, we essentially have that the probability of finding exceptional Italian teen agers not loving sport is higher than the one of finding exceptional ones not using apps, but both are typical properties of an Italian teen ager. In case the ontology engineer needs to formalize that typical Italian teen agers do not love sport, he just need to have $\mathbf{T}(\mathit{ItalianTeenAger}) \sqsubseteq_p \neg \mathit{SportLover}$ in his KB with a suitable p.

4 Extensions of ABox

Given a KB, we define the finite set $\mathfrak{Tip}$ of concepts occurring in the scope of the typicality operator, i.e. $\mathfrak{Tip} = \{C \mid \mathbf{T}(C) \sqsubseteq_p D \in \text{KB}\}$. Given an individual a explicitly named in the ABox, we define the set of typicality assumptions $\mathbf{T}(C)(a)$ that can be minimally entailed from KB in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathit{RaCl}}$, with $C \in \mathfrak{Tip}$. We then consider an ordered set $\mathfrak{Tip}_{\mathcal{A}}$ of pairs (a, C) of all possible assumptions $\mathbf{T}(C)(a)$, for all concepts $C \in \mathfrak{Tip}$ and all individual constants a in the ABox.

Definition 5 **(Assumptions in $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$).** *Given an $\mathcal{ALC} + \mathbf{T}^{P}_{\mathbf{R}}$ KB=($\mathcal{T}, \mathcal{A}$), let $\mathcal{T}'$ be the set of inclusions of $\mathcal{T}$ without probabilities, namely*

$$\mathcal{T}' = \{\mathbf{T}(C) \sqsubseteq D \mid \mathbf{T}(C) \sqsubseteq_p D \in \mathcal{T}\} \cup \{C \sqsubseteq D \in \mathcal{T}\}.$$

Given a finite set of concepts $\mathfrak{Tip}$, we define, for each individual name a occurring in $\mathcal{A}$: $\mathfrak{Tip}_a = \{C \in \mathfrak{Tip} \mid (\mathcal{T}', \mathcal{A}) \models_{\mathcal{ALC}+\mathbf{T}^{RaCl}_{\mathbf{R}}} \mathbf{T}(C)(a)\}$. We also define $\mathfrak{Tip}_{\mathcal{A}} = \{(a, C) \mid C \in \mathfrak{Tip}_a$ and a occurs in $\mathcal{A}\}$ and we impose an order on its elements: $\mathfrak{Tip}_{\mathcal{A}} = [(a_1, C_1), (a_2, C_2), \ldots, (a_n, C_n)]$. Furthermore, we define the ordered multiset *$\mathcal{P}_{\mathcal{A}} = [p_1, p_2, \ldots, p_n]$, respecting the order imposed on $\mathfrak{Tip}_{\mathcal{A}}$, where*

$$p_i = \prod_{j=1}^{m} p_{ij} \textit{ for all } \mathbf{T}(C_i) \sqsubseteq_{p_{i1}} D_1, \mathbf{T}(C_i) \sqsubseteq_{p_{i2}} D_2, \ldots, \mathbf{T}(C_i) \sqsubseteq_{p_{im}} D_m \textit{ in } \mathcal{T}.$$

The ordered multiset $\mathcal{P}_{\mathcal{A}}$ is a tuple of the form $[p_1, p_2, \ldots, p_n]$, where p_i is the probability of the assumption $\mathbf{T}(C)(a)$, such that $(a, C) \in \mathfrak{Tip}_{\mathcal{A}}$ at position i. p_i is the product of all the probabilities p_{ij} of typicality inclusions $\mathbf{T}(C) \sqsubseteq_{p_{ij}} D$ in the TBox.

Following the basic idea underlying surprising scenarios outlined in [15], we consider different extensions $\widetilde{\mathcal{A}_i}$ of the ABox and we equip them with a probability $\mathbb{P}_i$. Starting from $\mathcal{P}_{\mathcal{A}} = [p_1, p_2, \ldots, p_n]$, the first step is to build all alternative tuples where 0 is used in place of some p_i to represent that the corresponding typicality assertion $\mathbf{T}(C)(a)$ is no longer assumed (Definition 6). Furthermore, we define the *extension* of the ABox corresponding to a string so obtained (Definition 7). In this way, the highest probability is assigned to the extension of the ABox corresponding to $\mathcal{P}_{\mathcal{A}}$, where all typicality assumptions are considered. The probability decreases in the other extensions, where some typicality assumptions are discarded, thus 0 is used in place of the corresponding p_i. The probability of an extension $\widetilde{\mathcal{A}_i}$ corresponding to a string $\mathcal{P}_{\mathcal{A}_i} = [p_{i1}, p_{i2}, \ldots, p_{in}]$ is defined as the product of probabilities p_{ij} when $p_{ij} \neq 0$, i.e. the probability of the corresponding typicality assumption when this is selected for the extension, and $1 - p_j$ when $p_{ij} = 0$, i.e. the corresponding typicality assumption is discarded, that is to say the extension contains an exception to the inclusion.

Definition 6 **(Strings of possible assumptions $\mathbb{S}$).** *Given a KB=($\mathcal{T}, \mathcal{A}$), let the set $\mathfrak{Tip}_{\mathcal{A}}$ and $\mathcal{P}_{\mathcal{A}} = [p_1, p_2, \ldots, p_n]$ be as in Definition 5. We define the set $\mathbb{S}$ of all the* strings of possible assumptions *with respect to KB as*

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

Definition 7 **(Extension of ABox).** *Let KB=($\mathcal{T}, \mathcal{A}$), $\mathcal{P}_{\mathcal{A}} = [p1, p_2, \ldots, p_n]$ and $\mathfrak{Tip}_{\mathcal{A}} = [(a_1, C_1), (a_2, C_2), \ldots, (a_n, C_n)]$ as in Definition 5. Given a string of possible assumptions $[s_1, s_2, \ldots, s_n] \in \mathbb{S}$ of Definition 6, we define the extension $\widetilde{\mathcal{A}}$ of $\mathcal{A}$ with respect to $\mathfrak{Tip}_{\mathcal{A}}$ and $\mathbb{S}$ as $\widetilde{\mathcal{A}} = \{\mathbf{T}(C_i)(a_i) \mid (a_i, C_i) \in \mathfrak{Tip}_{\mathcal{A}}$ and $s_i \neq 0\}$. We also define the probability of $\widetilde{\mathcal{A}}$ as $\mathbb{P}_{\widetilde{\mathcal{A}}} = \prod_{i=1}^{n} \chi_i$ where $\chi_i = \begin{cases} p_i & \textit{if } s_i \neq 0 \\ 1 - p_i & \textit{if } s_i = 0 \end{cases}$*

It can be observed that, in $\mathcal{ALC} + \mathbf{T}^{RaCl}_{\mathbf{R}}$, the set of typicality assumptions that can be inferred from a KB corresponds to the extension of $\mathcal{A}$ corresponding to the string

$\mathcal{P}_{\mathcal{A}}$ (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 $\mathcal{ALC} + \mathbf{T_R}$, no typicality assumptions 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. It is easy to observe that we obtain a probability distribution over extensions of $\mathcal{A}$.

Example 2. Given a KB=$(\mathcal{T}, \mathcal{A})$, let the only typicality inclusions in $\mathcal{T}$ be $\{\mathbf{T}(C) \sqsubseteq_{0.6} D, \mathbf{T}(E) \sqsubseteq_{0.85} F\}$. Let a and b be the only individual constants occurring in $\mathcal{A}$. Suppose also that $\mathbf{T}(C)(a)$, $\mathbf{T}(C)(b)$, and $\mathbf{T}(E)(b)$ are entailed from KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$. We have that $\mathfrak{Tip}_{\mathcal{A}} = \{(a, C), (b, C), (b, E)\}$ and $\mathcal{P}_{\mathcal{A}} = [0.6, 0.6, 0.85]$. All possible strings, corresponding extensions of $\mathcal{A}$ and probabilities are shown in Table 1.

String	**Extension**	**Probability**
[0.6, 0.6, 0.85]	$\widetilde{\mathcal{A}_1} = \{\mathbf{T}(C)(a), \mathbf{T}(C)(b), \mathbf{T}(E)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_1}} = 0.6 \times 0.6 \times 0.85 = 0.306$
[0, 0, 0.85]	$\widetilde{\mathcal{A}_2} = \{\mathbf{T}(E)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_2}} = (1 - 0.6) \times (1 - 0.6) \times 0.85 = 0.136$
[0, 0.6, 0]	$\widetilde{\mathcal{A}_3} = \{\mathbf{T}(C)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_3}} = (1 - 0.6) \times 0.6 \times (1 - 0.85) = 0.036$
[0.6, 0, 0]	$\widetilde{\mathcal{A}_4} = \{\mathbf{T}(C)(a)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_4}} = 0.6 \times (1 - 0.6) \times (1 - 0.85) = 0.036$
[0, 0.6, 0.85]	$\widetilde{\mathcal{A}_5} = \{\mathbf{T}(C)(b), \mathbf{T}(E)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_5}} = (1 - 0.6) \times 0.6 \times 0.85 = 0.204$
[0.6, 0, 0.85]	$\widetilde{\mathcal{A}_6} = \{\mathbf{T}(C)(a), \mathbf{T}(E)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_6}} = 0.6 \times (1 - 0.6) \times 0.85 = 0.204$
[0.6, 0.6, 0]	$\widetilde{\mathcal{A}_7} = \{\mathbf{T}(C)(a), \mathbf{T}(C)(b)\}$	$\mathbb{P}_{\widetilde{\mathcal{A}_7}} = 0.6 \times 0.6 \times (1 - 0.85) = 0.054$
[0, 0, 0]	$\widetilde{\mathcal{A}_8} = \emptyset$	$\mathbb{P}_{\widetilde{\mathcal{A}_8}} = (1{-}0.6) \times (1{-}0.6) \times (1{-}0.85) = 0.024$
	$\mathbb{P}_{\widetilde{\mathcal{A}_1}} + \mathbb{P}_{\widetilde{\mathcal{A}_2}} + \ldots + \mathbb{P}_{\widetilde{\mathcal{A}_8}} =$	1

Table 1. Plausible extensions of the ABox of Example 2.

5 A Decision Procedure for Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$

Given KB and a query F, we distinguish two cases:

1. if F is an inclusion $C \sqsubseteq D$, then it is entailed from KB if it is minimally entailed from KB' in the nonmonotonic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, where KB' is obtained from KB by removing probabilities of exceptions, i.e. by replacing each typicality inclusion $\mathbf{T}(C) \sqsubseteq_p D$ with $\mathbf{T}(C) \sqsubseteq D$;
2. if F is an ABox fact $C(a)$, then it is entailed from KB if it is entailed in the monotonic $\mathcal{ALC} + \mathbf{T_R}$ from the knowledge bases including the extensions of the ABox of Definition 7. More in detail, we provide both (i) a notion of entailment restricted to scenarios whose probabilities belong to a given range and (ii), similarly to [13], a notion of probability of the entailment of a query $C(a)$, as the sum of the probabilities of all extensions from which $C(a)$ is so entailed.

Here below are the formal definition of entailment of a query in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$. We distinguish the case in which the query is a TBox inclusion from the one in which it is an ABox assertion. Notice that, in the former case, probabilities do not play any role, and, as already mentioned, entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ corresponds to entailment in the nonmonotonic Description Logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$.

Definition 8 (Entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$). *Given a KB=($\mathcal{T}, \mathcal{A}$), given $\mathfrak{Tip}$ a set of concepts, two real numbers $p, q \in (0, 1]$, let $\mathcal{E} = \{\widetilde{A_1}, \widetilde{A_2}, \ldots, \widetilde{A_k}\}$ be the set of extensions of $\mathcal{A}$ of Definition 7 with respect to $\mathfrak{Tip}$, whose probabilities are such that $p \leq \mathbb{P}_1 \leq q, p \leq \mathbb{P}_2 \leq q, \ldots, p \leq \mathbb{P}_k \leq q$. Let $\mathcal{T}' = \{\mathbf{T}(C) \sqsubseteq D \mid \mathbf{T}(C) \sqsubseteq_r D \in \mathcal{T}\} \cup \{C \sqsubseteq D \in \mathcal{T}\}$. Given a query F, we say that F is entailed from KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ in range $\langle p, q\rangle$, written KB $\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$, as follows:*

- *if F is either an inclusion $C \sqsubseteq D$ or a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$, if ($\mathcal{T}', \mathcal{A}$) $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} F$;*
- *if F is an ABox assertion $C(a)$, where $a \in \mathcal{O}$, if ($\mathcal{T}', \mathcal{A} \cup \widetilde{\mathcal{A}_i}$) $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}} F$ for all $\widetilde{\mathcal{A}_i} \in \mathcal{E}$. We also define the probability of the entailment of a query as $\mathbb{P}(F) = \sum_{i=1}^{k} \mathbb{P}_i$.*

Our decision procedure checks whether a query F is entailed from a given KB as in Definition 8. We then exploit such decision procedure to show that the problem of entailment in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ is in EXPTIME. This allows us to conclude that reasoning about typicality and defeasible inheritance with probabilities of exceptions 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 [1].

Given an $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ KB=($\mathcal{T}, \mathcal{A}$) and a query F, we define a procedure computing the following four steps:

1. compute the set $\mathfrak{Tip}_a$ of all typicality assumptions that are minimally entailed from the knowledge base in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$;
2. compute all possible $\widetilde{A_i}$ extensions of the ABox and compute their probabilities;
3. select the extensions whose probabilities belong to a given range $\langle p, q\rangle$;
4. check whether the query F is entailed from all the selected extensions in the monotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$.

Step 4 is based on reasoning in the monotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$: to this aim, the procedure relies on a polynomial encoding of $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ into $\mathcal{ALC}$ introduced in [16]. Step 1 is based on reasoning in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$: in this case, the procedure computes the rational closure of an $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ knowledge base by means of the algorithm introduced in [11, 17]. Also the algorithm computing the rational closure relies on reasoning in the monotonic logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$, then on the above mentioned polynomial encoding in $\mathcal{ALC}$.

Let KB=($\mathcal{T}, \mathcal{A}$) be an $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ knowledge base. Let $\mathcal{T}'$ be the set of inclusions of $\mathcal{T}$ without probabilities of exceptions: $\mathcal{T}' = \{\mathbf{T}(C) \sqsubseteq D \mid \mathbf{T}(C) \sqsubseteq_r D \in \mathcal{T}\} \cup \{C \sqsubseteq D \in \mathcal{T}\}$, that the procedure will consider in order to reason in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$ and $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ for checking query entailment and finding all plausible typicality assumptions, respectively. Other inputs of the procedure are the finite set of concepts

$\mathfrak{Tip}$, a query F, and two real numbers $p, q \in (0, 1]$ describing a range of probabilities. If F is $C \sqsubseteq D$ (where C could be $\mathbf{T}(C')$), we just need to check whether $(\mathcal{T}', \mathcal{A})$ $\models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} C \sqsubseteq D$ in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$. If F is an ABox formula $C(a)$, Algorithm 1 builds all possible scenarios, computes their probabilities and then checks whether KB $\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$ if F holds in all those scenarios having a probability between p and q.

Algorithm 1 Entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$: KB $\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$

---

procedure ENTAILMENT($(\mathcal{T}, \mathcal{A}), \mathcal{T}', F, \mathfrak{Tip}, p, q$)
    if F is of the form $C \sqsubseteq D$ **then** ▷ *If F is an inclusion, rely on $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$*
        return $(\mathcal{T}', \mathcal{A}) \models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} F$
▷ *Otherwise, F is an ABox assertion of the form $C(a)$*
    $\mathfrak{Tip}_{\mathcal{A}} \leftarrow \emptyset$ ▷ build the set $\mathbb{S}$ of possible assumptions
    for each $C \in \mathfrak{Tip}$ **do**
        for each individual $a \in \mathcal{A}$ **do** ▷ Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$
            if $(\mathcal{T}', \mathcal{A}) \models_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}} \mathbf{T}(C)(a)$ **then** $\mathfrak{Tip}_{\mathcal{A}} \leftarrow \mathfrak{Tip}_{\mathcal{A}} \cup \{\mathbf{T}(C)(a)\}$
    $\mathcal{P}_{\mathcal{A}} \leftarrow \emptyset$ ▷ compute the probabilities of Definition 5 given $\mathcal{T}$ and $\mathfrak{Tip}_{\mathcal{A}}$
    for each $C \in \mathfrak{Tip}$ **do**
        $\Pi_C \leftarrow 1$
        for each $\mathbf{T}(C) \sqsubseteq_p D \in \mathcal{T}$ **do** $\Pi_C \leftarrow \Pi_C \times p$
        $\mathcal{P}_{\mathcal{A}} \leftarrow \mathcal{P}_{\mathcal{A}} \cup \Pi_C$
    $\mathbb{S} \leftarrow$ *build strings of possible assumptions as in Definition 6 given $\mathfrak{Tip}_{\mathcal{A}}$ and $\mathcal{P}_{\mathcal{A}}$*
    $\mathcal{E} \leftarrow \emptyset$ ▷ build extensions of $\mathcal{A}$
    for each $s_i \in \mathbb{S}$ **do**
        build the extension $\widetilde{\mathcal{A}_i}$ corresponding to s_i and compute $\mathbb{P}_{\widetilde{\mathcal{A}_i}}$ as in Definition 7
        if $p \leq \mathbb{P}_{\widetilde{\mathcal{A}_i}} \leq q$ **then** $\mathcal{E} \leftarrow \mathcal{E} \cup \widetilde{\mathcal{A}_i}$ ▷ select extensions with probability in $\langle p, q \rangle$
    for each $\widetilde{\mathcal{A}_i} \in \mathcal{E}$ **do** ▷ query entailment in $\mathcal{ALC} + \mathbf{T_R}$
        if $(\mathcal{T}', \mathcal{A} \cup \widetilde{\mathcal{A}_i}) \not\models_{\mathcal{ALC}+\mathbf{T_R}} F$ **then return** KB $\not\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$
    return KB $\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$ ▷ F is entailed in all extensions

---

We can exploit the procedure of Algorithm 1 to show that the problem of entailment in the logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ is EXPTIME complete. This allows us to conclude that reasoning about typicality and defeasible inheritance with probabilities of exceptions is essentially inexpensive, since reasoning retains the same complexity class of the underlying standard $\mathcal{ALC}$, which is known to be EXPTIME-complete [1].

Theorem 1 (Complexity of entailment). *Given a KB in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, real numbers $p, q \in (0, 1]$ and a query F whose size is polynomial in the size of KB, the problem of checking whether KB $\models^{\langle p,q \rangle}_{\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}} F$ is* EXPTIME*-complete.*

Proof. (*sketch*, [14]) The algorithm checks, for each concept $C \in \mathfrak{Tip}$ and for each individual a whether $\mathbf{T}(C)(a)$ is minimally entailed from KB in nonmonotonic $\mathcal{ALC}+$

$\mathbf{T}_{\mathbf{R}}^{RaCl}$. Let n be the length of the string representing KB. By definition, the size of $\mathfrak{Tip}$ is $O(n)$. There are $O(n^2)$ concepts $\mathbf{T}(C)(a)$, for each of them the algorithm relies on reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$, which is in EXPTIME [11]. Building $\mathcal{P}_{\mathcal{A}}$ can be solved with $O(n^2)$ operations. The algorithm considers all possible strings obtained by assuming (or not) each typicality assumption $\mathbf{T}(C)(a)$ (they are $O(n^2)$) in order to compute the set $\mathbb{S}$ of plausible extensions. For each s_i, we have either $s_i = 0$ or $s_i \neq 0$, then $2 \times 2 \times \ldots \times 2$ different strings, thus $\mathbb{S}$ has exponential size in n. Selecting extensions whose probabilities $\mathbb{P}_{\widetilde{\mathcal{A}_i}}$ are in the range $[p, q]$ can be solved in EXPTIME, then the algorithm relies on reasoning in monotonic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$, in order to check whether F is entailed in selected extensions in $\mathcal{E}$, whose size is $O(2^n)$: we have $O(2^n)$ calls to query entailment in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$, which is EXPTIME-complete. □

Let us now conclude Example 1 introduced in Section 3 in the light of the definitions provided above.

Example 3 (Reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ part 2). Suppose that the ABox is

$$\mathcal{A} = \{AtypicalDepressed(john), ProstateCancerPatient(greg)\},$$

we can consider two typicality assumptions: (a) $\mathbf{T}(AtypicalDepressed)(john)$ and (b) $\mathbf{T}(ProstateCancerPatient)(greg)$. We distinguish among four different extensions:

1. both (a) and (b) are assumed: in this scenario, whose probability is $0.7 \times (0.6 \times 0.8) = 0.336$, we conclude that both John and Greg have mood swings, and that Greg has nocturia;
2. we assume (b) but not (a): this scenario has probability $(1 - 0.7) \times (0.6 \times 0.8) = 0.144$, and we can only conclude $\exists hasSymptom.MoodReactivity(greg)$ and $\exists hasSymptom.Nocturia(greg)$;
3. we assume (a) and not (b): this scenario, having a probability $0.7 \times (1 - (0.6 \times 0.8)) = 0.364$, allows us to conclude $\exists hasSymptom.MoodReactivity(john)$;
4. neither (a) nor (b) is added to $\mathcal{A}$: here the probability is $(1 - 0.7) \times (1 - (0.6 \times 0.8)) = 0.156$, but we are not able to conclude anything about John and Greg.

The probability that John has mood swings is defined as the sum of the probabilities of scenarios where such inference can be performed, namely scenarios (1) and (3), and it is therefore $0.336 + 0.364 = 0.7$. Similarly, the probability that Greg has nocturia and mood swings is $0.336 + 0.144 = 0.48$.

6 Conclusions

We have described the Description Logic $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ introduced in [14], which extends the nonmonotonic Description Logic of typicality $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{RaCl}$ by means of probabilities equipping typicality inclusions. Probabilities of exceptions are then used in order to reason about plausible scenarios, obtained by selecting only some – i.e., not necessarily all – typicality assumptions and whose probabilities belong to a given and fixed range. We have also presented a decision procedure for reasoning in $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, and we have shown that such a procedure can be exploited in order to estimate the complexity of the proposed logic, namely to show that reasoning in DLs with rational closure and probabilities of exceptions remains in EXPTIME as in the underlying standard $\mathcal{ALC}$.

The logic $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$, as well as the underlying $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}$, are based on the rational closure, then they inherit its virtues, but also its weakness. It is well known that the main advantage of the rational closure is related to its good computational properties. However, rational closure is affected by the "all or nothing" behavior, in the sense that it does not allow one to separately reason about the inheritance of different properties. For instance, let us again recall the example in the Introduction: we have that typical Italian teen agers make use of apps for mobile devices, and normally they also listen to trap music. Furthermore, we can consistently express that, normally, convict Italian teen agers do not make use of apps. As a consequence, they are recognized as untypical Italian teen agers, then no inheritance of typical properties is possible, for instance it is not possible to infer that they listen to trap music. The problem also affects the definition of scenarios in the logic $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$: if $\mathbf{T}(\mathit{ItalianTeenAger})(\mathit{simone})$ is a typicality assumption to be considered in the construction of different scenarios (since it is entailed in $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{RaCl}$ from the knowledge base), then Simone inherits all the properties of typical Italian teen agers. On the contrary, if $\mathbf{T}(\mathit{ConvictItalianTeenAger})(\mathit{simone})$ is the typicality assertion to be considered for the scenarios generation, no inheritance of typical Italian teen agers is possible for Simone. In order to solve this problem, a strengthening of a rational closure-like algorithm with defeasible inheritance networks has been studied by [18]. In [19] the author has proposed an alternative semantics by considering models equipped with multiple preference relations, whence with multiple "typicality" operators. In this variant, it should be possible to distinguish different aspects of typicality/exceptionality and consequently to avoid the "all or nothing" behavior of rational closure with respect to property inheritance.

Several nonmonotonic extensions of DLs have been proposed in the literature in order to reason about inheritance with exceptions, essentially based on the integration of DLs with well established nonmonotonic reasoning mechanisms [2–4, 20, 5, 6, 8], ranging from Reiter's defaults to minimal knowledge and negation as failure (see [9, 8] for a detailed presentation). In none of them, probability of exceptions in concept inclusions is taken into account, as far as we know.

Probabilistic extensions of DLs, allowing one to label inclusions (and facts) with degrees representing probabilities, have been introduced in [12, 13]. In this approach, called DISPONTE, the authors propose the integration of probabilistic information with DLs based on the distribution semantics for probabilistic logic programs [21]. The basic idea is to label inclusions of the TBox as well as facts of the ABox with a real number between 0 and 1, representing their probabilities, assuming that each axiom is independent from each others. The resulting knowledge base defines a probability distribution over *worlds*: roughly speaking, a world is obtained by choosing, for each axiom of the KB, whether it is considered as true of false. The distribution is further extended to queries and the probability of the entailment of a query is obtained by marginalizing the joint distribution of the query and the worlds. There are two main differences between the logic $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ proposed in this work and probabilistic DLs. On the one hand, as already mentioned in the Introduction, in the logic $\mathcal{ALC}+\mathbf{T}_{\mathbf{R}}^{\mathsf{P}}$ probabilities are used in order to express different degrees of admissibility of exceptions with respect to such typicality inclusions. Probabilities are then the basis of different scenarios built by assuming – or not – that individuals are typical instances of a given concept. On the contrary, in DISPONTE probabilities are used to capture a notion of uncertainty about information of the KB, therefore an inclusion $C \sqsubseteq D$ having a very low probability p has a significantly different meaning with respect to an inclusion $\mathbf{T}(C) \sqsubseteq_p D$, representing anyway a typical property: normally, Cs are Ds, even if with a high probability of having exceptions to such typical inclusion. On the other hand, in $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ probabilities are restricted to typicality inclusions only. On the contrary, in DISPONTE probabilities can be associated to concept inclusions as well as to ABox facts.

In [15] a nonmonotonic procedure for reasoning about *surprising* scenarios in DLs has been proposed. In this approach, the DL $\mathcal{ALC} + \mathbf{T_R}$ is extended by inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$, where d is a *degree of expectedness*. In this logic, cardinality restrictions play a fundamental role in order to "filter" extended ABoxes and entailment is restricted to minimal scenarios, determined by a partial order among such extended ABoxes, whereas in our logic $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$, entailment is defined in terms of the probability of a given scenario and can be used to estimate the probability of a given query.

In future work we aim at extending the logic $\mathcal{ALC} + \mathbf{T}^{\mathsf{P}}_{\mathbf{R}}$ to more expressive Description Logics, such as those underlying the standard language for ontology engineering OWL. As a first step, in [16] the logic with the typicality operator and the rational closure construction have been applied to the logic $\mathcal{SHIQ}$.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: The Description Logic Handbook - Theory, Implementation, and Applications, 2nd edition. Cambridge (2007)
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. 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
5. 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
6. Casini, G., Straccia, U.: Defeasible Inheritance-Based Description Logics. Journal of Artificial Intelligence Research (JAIR) **48** (2013) 415–473
7. 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
8. Bonatti, P.A., Faella, M., Petrova, I., Sauro, L.: A new semantics for overriding in description logics. Artificial Intelligence **222** (2015) 1–48
9. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: ALC+T: a preferential extension of description logics. Fundamenta Informaticae **96** (2009) 341–372
10. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**(1) (1992) 1–60

11. 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
12. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Reasoning with probabilistic ontologies. In Yang, Q., Wooldridge, M., eds.: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015, AAAI Press (2015) 4310–4316
13. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Probabilistic description logics under the distribution semantics. Semantic Web **6**(5) (2015) 477–501
14. Pozzato, G.L.: Reasoning in description logics with typicalities and probabilities of exceptions. In Antonucci, A., Cholvy, L., Papini, O., eds.: Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 14th European Conference, ECSQARU 2017, Lugano, Switzerland, July 10-14, 2017, Proceedings. Volume 10369 of Lecture Notes in Computer Science., Springer (2017) 409–420
15. Pozzato, G.L.: Reasoning about surprising scenarios in description logics of typicality. In Adorni, G., Cagnoni, S., Gori, M., Maratea, M., eds.: Advances in Artificial Intelligence: Proceedings of the 15th International Conference of the Italian Association for Artificial Intelligence. Volume 10037 of Lecture Notes in Artificial Intelligence LNAI., Genova, Springer (November 2016) 1–15
16. 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
17. Giordano, L., Gliozzi, V., Pozzato, G.L., Renzulli, R.: An efficient reasoner for description logics of typicality and rational closure. In Artale, A., Glimm, B., Kontchakov, R., eds.: Proceedings of the 30th International Workshop on Description Logics, Montpellier, France, July 18-21, 2017. Volume 1879 of CEUR Workshop Proceedings., CEUR-WS.org (2017)
18. Casini, G., Straccia, U.: Defeasible Inheritance-Based Description Logics. In Walsh, T., ed.: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), Barcelona, Spain, Morgan Kaufmann (July 2011) 813–818
19. Gliozzi, V.: Reasoning about Multiple Aspects in Rational Closure for DLs. In Adorni, G., Cagnoni, S., Gori, M., Maratea, M., eds.: Advances in Artificial Intelligence: Proceedings of the 15th International Conference of the Italian Association for Artificial Intelligence. Volume 10037 of Lecture Notes in Artificial Intelligence LNAI., Genova, Springer (November 2016) 392–405
20. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: A NonMonotonic Description Logic for Reasoning About Typicality. Artificial Intelligence **195** (2013) 165 – 202
21. Sato, T.: A statistical learning method for logic programs with distribution semantics. In Sterling, L., ed.: Logic Programming, Proc. of the 12th International Conference on Logic Programming, MIT Press (1995) 715–729