IRIS AperTO

UNIVERSITÀ DEGLI STUDI DI TORINO

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics

This is the author's manuscript

Original Citation:

Availability:

This version is available http://hdl.handle.net/2318/1736642 since 2021-08-01T23:57:08Z

Published version:

DOI:10.1080/0952813X.2019.1672799

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

09 November 2026

A Description Logic Framework for Commonsense Conceptual Combination Integrating Typicality, Probabilities and Cognitive Heuristics

Antonio Lieto[a] and Gian Luca Pozzato[b]

[a]Dipartimento di Informatica, Università di Torino, Italy and ICAR-CNR, Istituto di Calcolo e Reti ad Alte Prestazioni, Palermo, Italy;
[b]Dipartimento di Informatica, Università di Torino, Italy

ARTICLE HISTORY
Compiled August 17, 2019

ABSTRACT
We propose a nonmonotonic Description Logic of typicality able to account for the phenomenon of the combination of prototypical concepts. The proposed logic relies on the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$, whose semantics is based on the notion of rational closure, as well as on the distributed semantics of probabilistic Description Logics, and is equipped with a cognitive heuristic used by humans for concept composition.

We first extend the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ by typicality inclusions of the form $p :: \mathbf{T}(C) \sqsubseteq D$, whose intuitive meaning is that "we believe with degree p about the fact that typical Cs are Ds". As in the distributed semantics, we define different scenarios containing only some typicality inclusions, each one having a suitable probability. We then exploit such scenarios in order to ascribe typical properties to a concept C obtained as the combination of two prototypical concepts. We also show that reasoning in the proposed Description Logic is EXPTIME-complete as for the underlying standard Description Logic $\mathcal{ALC}$.

KEYWORDS
Nonmonotonic reasoning; description logics; common-sense reasoning; cognitive modelling.

1. Introduction

Inventing novel concepts by combining the typical knowledge of pre-existing ones is among the most creative cognitive abilities exhibited by humans. This generative phenomenon highlights some crucial aspects of the knowledge processing capabilities in human cognition and concerns high-level capacities associated to creative thinking and problem solving. Still, it represents an open challenge in the field od Artificial Intelligence (AI) [Boden (1998); Colton et al. (2012)]. Dealing with this problem requires, from an AI perspective, the harmonization of two conflicting requirements that are hardly accommodated in symbolic systems (including formal ontologies [Frixione and Lieto (2011)]): the need of a syntactic and semantic compositionality (typical of logical systems) and that one concerning the exhibition of typicality effects. According to a well-known argument, in fact, prototypes are not compositional [Fodor (1981); Gleitman et al. (2012); Murphy (2002); Osherson and Smith (1981)]. The argument runs as follows: consider a concept like *pet fish*. It results from the composition of

CONTACT Antonio Lieto. Email: antonio.lieto@unito.it
CONTACT Gian Luca Pozzato. Email: gianluca.pozzato@unito.it

the concept *pet* and of the concept *fish*. However, the prototype of *pet fish* cannot result from the composition of the prototypes of a pet and a fish: e.g. a typical pet is furry and warm, a typical fish is grayish, but a typical pet fish is neither furry and warm nor grayish (typically, it is red).

In this work, we provide a framework able to account for this type of human-like concept combination by explicitly relying on a formalization of the prototype theory [1]. We propose a nonmonotonic Description Logic (from now on DL) of typicality called $\mathbf{T}^{\mathsf{CL}}$ (typicality-based compositional logic). This logic combines two main ingredients. The first one relies on the DL of typicality $\mathcal{ALC} + \mathbf{T_R}$ introduced in [Giordano et al. (2015)]. In this logic, "typical" 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". As a difference with standard DLs, in the logic $\mathcal{ALC} + \mathbf{T_R}$ one can consistently express exceptions and reason about defeasible inheritance as well. For instance, a knowledge base can consistently express that "normally, athletes are in fit", whereas "sumo wrestlers usually are not in fit" by the typicality inclusions

$\mathbf{T}(Athlete) \sqsubseteq Fit$
$\mathbf{T}(SumoWrestler) \sqsubseteq \neg Fit,$

given that $SumoWrestler \sqsubseteq Athlete$. The semantics of the $\mathbf{T}$ operator is characterized by the properties of *rational logic* [Lehmann and Magidor (1992)], recognized as the core properties of nonmonotonic reasoning. $\mathcal{ALC} + \mathbf{T_R}$ is characterized by a minimal model semantics corresponding to an extension to DLs of a notion of *rational closure* as defined in [Lehmann and Magidor (1992)] for propositional logic: the idea is to adopt a preference relation among $\mathcal{ALC} + \mathbf{T_R}$ models, where intuitively a model is preferred to another one if it contains less exceptional elements, as well as a notion of *minimal entailment* restricted to models that are minimal with respect to such preference relation. As a consequence, $\mathbf{T}$ inherits well-established properties like *specificity* and *irrelevance*: in the example, the logic $\mathcal{ALC} + \mathbf{T_R}$ allows us to infer $\mathbf{T}(Athlete \sqcap Bald) \sqsubseteq Fit$ (being bald is irrelevant with respect to being in fit) and, if one knows that Hiroyuki is a typical sumo wrestler, to infer that he is not in fit, giving preference to the most specific information.

As a second ingredient, we rely on a distributed semantics similar to the DISPONTE semantics proposed by [Riguzzi et al. (2015a,b)] for probabilistic extensions of DLs, more precisely we extend the typicality inclusions of $\mathcal{ALC} + \mathbf{T_R}$ with probabilities intended as *degrees of belief*. As in the DISPONTE semantics, we make use of these degrees in order to define different scenarios, corresponding in our case to subsets of typicality inclusions. More in detail, we restrict our adoption of DISPONTE to label inclusions $\mathbf{T}(C) \sqsubseteq D$ with a real number between 0.5 and 1, representing the degree of belief in such an inclusion[2], assuming that each axiom is independent from each other (as in the DISPONTE semantics). The resulting knowledge base defines a probability distribution over *scenarios*: roughly speaking, a scenario is obtained by choosing, for each typicality inclusion, whether it is considered as true or false. As a consequence, two types of vagueness are expressed by a typicality inclusion $p :: \mathbf{T}(C) \sqsubseteq D$: on the one hand, the typicality operator $\mathbf{T}$ allows one to express that only typical members of C are also members of D, on the other hand, p allows one to have typical properties having different degrees of belief in each typicality inclusions. In a slight extension of the above example, we could have the need of representing that both the typical-

[1] Thus providing additional arguments in favor of the theoretical positions according to which it is possible to reconcile compositionality and prototypes [Hampton (2000); Prinz (2012)].

[2] We want to stress that, as in any probabilistic formal framework, probabilities are assumed to come from an application domain. This is true also for other frameworks such as, for example, fuzzy logics or probabilistic extensions of logic programs. In this paper, we focus on the proposal of the formalism itself, therefore the machinery for obtaining probabilities from a dataset of the application domain is out of the scope.

ity inclusions about athletes and sumo wrestlers have a probability of 80%, whereas we also believe that athletes are usually young with a higher probability of 95%, with the following KB:

(1) $SumoWrestler \sqsubseteq Athlete$
(2) $0.8 :: \mathbf{T}(Athlete) \sqsubseteq Fit$
(3) $0.8 :: \mathbf{T}(SumoWrestler) \sqsubseteq \neg Fit$
(4) $0.95 :: \mathbf{T}(Athlete) \sqsubseteq YoungPerson$

We consider eight different scenarios, representing all possible combinations of typicality inclusion: as an example, we can consider the one in which (2) and (4) hold, whereas (3) does not. In this scenario, we have, for instance, that also sumo wrestlers are fit. We equip each scenario with a degree depending on those of the involved typicality inclusions. In the above example, the scenario containing only (2) and (4) has a probability $15,2\%$, obtained by the product of 0.8 (from 2), $1 - 0.8$ (since we do not include 3) and 0.95 (from 4).

As an additional element of the proposed formalization we employ a method inspired by cognitive semantics [Hampton (1987, 1988); Kamp and Partee (1995); Smith et al. (1988)] for the identification of a dominance effect between the concepts to be combined. Namely, for every combination, we distinguish a HEAD and a MODIFIER, where the HEAD represents the stronger element of the combination. The basic idea is as follows: given a KB and two concepts C_H and C_M occurring in it, where C_H is the HEAD and C_M the MODIFIER, we consider only *some* scenarios in order to define a revised knowledge base, enriched by typical properties of the combined concept $C \sqsubseteq C_H \sqcap C_M$. Such scenarios are those (i) consistent with respect to the initial knowledge base, (ii) not *trivial*, i.e. we discard those with the highest probability, containing either *all* properties that can be consistently ascribed to C or *all* properties of the HEAD that can be consistently ascribed to C, and (iii) giving preference to the typical properties of the HEAD C_H (with respect to those of C_M) having the highest probability.

We are able to exploit the logic $\mathbf{T}^{\mathsf{CL}}$ in two different perspectives. On the one hand, we show that it is able to capture well established examples in the literature of cognitive science concerning concept combination and, as such, we argue that $\mathbf{T}^{\mathsf{CL}}$ is a promising candidate to tackle the problem of typicality-based concept combination (Sections 4.1, 4.2 and 4.3). On the other hand, we use $\mathbf{T}^{\mathsf{CL}}$ as a tool for the generation and the exploration of novel creative concepts (Section 5), that could be useful in many applicative scenarios, ranging from video games to the creation of new movie or story characters.

As a further result, we show that the proposed approach is essentially inexpensive, in the sense that reasoning in $\mathbf{T}^{\mathsf{CL}}$ is EXPTIME-complete as for the underlying standard $\mathcal{ALC}$ Description Logic.

The plan of the paper is as follows. In Section 2 we recall the two semantics being the starting points of our proposal, namely the rational closure for the DL of typicality $\mathcal{ALC}+\mathbf{T_R}$ and the DISPONTE semantics for probabilistic DLs. In Section 3 we present the logic $\mathbf{T}^{\mathsf{CL}}$ for concept combination, and we show its reasoning complexity. In Section 4 we show that the proposed logic $\mathbf{T}^{\mathsf{CL}}$ is able to capture some well known and paradigmatic examples of concept combination coming from the cognitive science literature. In Section 5 we exploit the logic $\mathbf{T}^{\mathsf{CL}}$ in the application domain of computational creativity: i.e. we show how $\mathbf{T}^{\mathsf{CL}}$ can be used for inventing/generating novel concepts as the result of the combination of two (or more) prototypes. In Section 6 we proceed further by showing that the logic $\mathbf{T}^{\mathsf{CL}}$ can be iteratively applied to combine prototypical concepts already resulting from the combination of prototypes. We conclude in Section 7 by mentioning some related approaches addressing the problem of common-sense concept combination, as well as by discussing on possible

future works.

2. Background: DLs of typicality and probabilistic DLs

As introduced above, the main aim of this work is to introduce a nonmonotonic Description Logic able to deal with the combination of prototypical concepts. In order to achieve this goal, we exploit two well established logical frameworks:

- the nonmonotonic logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ based on a notion of rational closure for DLs.
- the DISPONTE semantics of probabilistic extensions of DLs

In this section we briefly recall such ingredients as well as Description Logics themselves, before introducing our proposal in Section 3.

2.1. Description Logics

The family of Description Logics is one of the most important formalisms of knowledge representation. DLs are reminiscent of the early semantic networks and of frame-based systems (for a detailed introduction see [Baader et al. (2003)]. They offer two key advantages: on the one hand, they have a well-defined semantics based on first-order logic and, on the other hand, they offer a good trade-off between expressivity of the language and computational complexity. DLs have been successfully implemented by a range of systems and they are at the base of languages for the semantic web such as OWL (Ontology Web Language).

A DL knowledge base (KB) comprises two components:

- the TBox, containing the definition of concepts (and possibly roles), and a specification of inclusions relations among them;
- the ABox containing instances of concepts and roles, in other words, properties and relations of individuals.

As an example, consider the following knowledge base, whose TBox contains the inclusion relations:

$$PhDStudent \sqsubseteq Student \tag{1}$$
$$Student \sqcap Worker \sqsubseteq TaxPayer \tag{2}$$
$$TaxPayer \sqsubseteq \exists hasId.(Code \sqcup NationalInsuranceNumber) \tag{3}$$

and whose ABox contains:

$$PhDStudent(max) \tag{4}$$
$$Worker(max) \tag{5}$$
$$NationalInsuranceNumber(12345) \tag{6}$$
$$hasId(max, 12345) \tag{7}$$

represents that PhD students are students (1), and that working students pay working taxes (2). Moreover, we have that tax payers possess an ID which is either a code or a NIN (National Insurance Number) (3). $Student$, $PhDStudent$, $Worker$, $Code$, $TaxPayer$, and $NationalInsuranceNumber$ are *concepts*. $hasId$ is a *role*, defining a relation between tax payers and IDs (codes and NINs). Furthermore, we have that Max is a both a student (4) and a worker (5), and he possesses the ID 12345 (7) which is a NIN (6).

The basic Description Logic is called $\mathcal{ALC}$. Concepts in $\mathcal{ALC}$ are built from an alphabet of concept names `C` and of role names `R`. Concept names define atomic concepts. Atomic

concepts and role names are then combined in order to build complex concepts: $C \sqcap D$ (intersection), $C \sqcup D$ (union), $\neg C$ (complement), $\forall R.C$ (universal qualified restriction) and $\exists R.C$ (existential qualified restriction). Extensions as well as fragments of $\mathcal{ALC}$ have been studied by considering other sets of operators.

Models in DLs are Kripke structures of the form

$$\langle \Delta^{\mathcal{I}}, .^{\mathcal{I}} \rangle$$

where $\Delta^{\mathcal{I}}$ is a non empty set of elements called the *domain*, whereas $.^{\mathcal{I}}$ is the *extension function* mapping each atomic concept C to the set of domain elements $C^{\mathcal{I}}$ being members of such concept, and mapping each atomic role R to pairs of domain elements $R^{\mathcal{I}}$ being related by such role. For instance, $Student^{\mathcal{I}}$ is the set of students of belonging to the domain under consideration. The function $.^{\mathcal{I}}$ is then extended to complex concepts in order to provide a suitable semantics to each connective/operator, that is to say:

- $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \backslash C^{\mathcal{I}}$, i.e. the extension of $\neg C$ is the complement of the extension of C, in other words it contains elements of $\Delta^{\mathcal{I}}$ not being C elements;
- $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, for instance the set of working students $Student \sqcap Worker$ corresponds to the intersection of the extension of $Student$ and the extension of $Worker$;
- $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$, for instance the extension of $Male \sqcup Female$ corresponds to the union of the extension of $Male$ and the extension of $Female$;
- $(\exists R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \exists (x, y) \in R^{\mathcal{I}}$ such that $y \in C^{\mathcal{I}}\}$, for instance the extension of $\exists tutoredBy.Professor$ contains all the elements of the domain having *at least* a professor as a tutor, i.e. those being related by role $tutoredBy$ to elements of the extension of $Professor$;
- $(\forall R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \forall (x, y) \in R^{\mathcal{I}}$ we have $y \in C^{\mathcal{I}}\}$, for instance the extension of $\forall hasChild.Female$ contains parents whose children are *all* female, in other words x belongs to $(\forall hasChild.Female)^{\mathcal{I}}$ if, for all pairs (x, y) in $hasChild^{\mathcal{I}}$ we have that y is a female (she belongs to the extension $Female^{\mathcal{I}}$).

One can reason about a DLs knowledge base by exploiting inference services like:

- satisfiability of the knowledge base: does the knowledge base admit a model?
- concept satisfiability: given a concept C, is there a model of the knowledge base assigning a non empty extension $C^{\mathcal{I}}$ to C?
- subsumption: given two concepts C and D, is C more general than D (i.e. $D^{\mathcal{I}} \subseteq C^{\mathcal{I}}$) in any model of the knowledge base?
- instance checking: given an individual a and a concept C, is a an instance of C (i.e. the domain element corresponding to a belongs to $C^{\mathcal{I}}$) in any model of the knowledge base?

In the above example, one can infer, for instance, that the knowledge base is consistent, whereas it would be not if we added the inclusion $(*)$ $Student \sqsubseteq \neg Worker$ (students are not workers, therefore there does not exist any model for the knowledge base, since Max is a working student). The concept $Worker \sqcap Student$ (working student) is satisfiable, whereas it would be not in presence of the inclusion $(*)$. Moreover, we can infer that Max is a tax payer

$$TaxPayer(max),$$

and that all working students have an ID

$$Student \sqcap Worker \sqsubseteq \exists hasId.(Code \sqcup NationalInsuranceNumber).$$

2.2. Reasoning about typicality in DLs: the nonmonotonic logic $\mathcal{ALC} + \mathbf{T_R}$

Since the very objective of the TBox is to build a taxonomy of concepts, the need of representing prototypical properties and of reasoning about defeasible inheritance of such properties easily arises. The traditional approach is to handle defeasible inheritance by integrating some kind of nonmonotonic reasoning mechanism: this has led to study nonmonotonic extensions of DLs [Baader and Hollunder (1995a); Bonatti et al. (2009); Donini et al. (2002a)] with the objective of fulfilling the following desiderata for such an extension:

- The (nonmonotonic) extension must have a clear semantics and should be based on the same semantics as the underlying monotonic DL.
- The extension should allow to specify prototypical properties in a natural and direct way.
- The extension must be decidable, if so is the underlying monotonic DL and, possibly, computationally effective.

A simple but powerful nonmonotonic extension of DLs is proposed in [Giordano et al. (2009, 2013a,b, 2015)]. In this approach, "typical" or "normal" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying DL. The logic so obtained is called $\mathcal{ALC} + \mathbf{T_R}$. 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 given by means of a set of postulates that are a reformulation of axioms and rules of nonmonotonic entailment in rational logic **R** [Lehmann and Magidor (1992)]: in this respect an assertion of the form $\mathbf{T}(C) \sqsubseteq D$ is equivalent to the conditional assertion $C \mathrel{|\!\sim} D$ in **R**. The basic ideas are as follows: given a domain $\Delta^{\mathcal{I}}$ and an evaluation function $.^{\mathcal{I}}$, one can define a function $f_{\mathbf{T}} : Pow(\Delta^{\mathcal{I}}) \longmapsto Pow(\Delta^{\mathcal{I}})$ that selects the *typical* instances of any $S \subseteq \Delta^{\mathcal{I}}$; in case $S = C^{\mathcal{I}}$ for a concept C, the selection function selects the typical instances of C, namely:

$$(\mathbf{T}(C))^{\mathcal{I}} = f_{\mathbf{T}}(C^{\mathcal{I}}).$$

$f_{\mathbf{T}}$ has the following properties for all subsets S of $\Delta^{\mathcal{I}}$, that are essentially a restatement of the properties characterizing rational logic **R**:

$(f_{\mathbf{T}} - 1)$ $f_{\mathbf{T}}(S) \subseteq S$
$(f_{\mathbf{T}} - 2)$ if $S \neq \emptyset$, then also $f_{\mathbf{T}}(S) \neq \emptyset$
$(f_{\mathbf{T}} - 3)$ if $f_{\mathbf{T}}(S) \subseteq R$, then $f_{\mathbf{T}}(S) = f_{\mathbf{T}}(S \cap R)$
$(f_{\mathbf{T}} - 4)$ $f_{\mathbf{T}}(\bigcup S_i) \subseteq \bigcup f_{\mathbf{T}}(S_i)$
$(f_{\mathbf{T}} - 5)$ $\bigcap f_{\mathbf{T}}(S_i) \subseteq f_{\mathbf{T}}(\bigcup S_i)$
$(f_{\mathbf{T}} - 6)$ if $f_{\mathbf{T}}(S) \cap R \neq \emptyset$, then $f_{\mathbf{T}}(S \cap R) \subseteq f_{\mathbf{T}}(S)$

The semantics of the $\mathbf{T}$ operator can be equivalently formulated in terms of *rational models* [Giordano et al. (2015)]: 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 $\Delta^{\mathcal{I}}$, if $x < y$ then either $x < z$ or $z < y$) relation over $\Delta^{\mathcal{I}}$. 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 $x \in \Delta^{\mathcal{I}}$ is a *typical instance* of some concept C if $x \in C^{\mathcal{I}}$ and there is no C-element in $\Delta^{\mathcal{I}}$ *more typical* than x. Elements in $\Delta^{\mathcal{I}}$ are then organized in different "levels" or "ranks" by the modularity of $<$, where each element x at rank i is incomparable with each other (i.e., for each y at rank i neither $x < y$ nor $y < x$) and it is more normal than all elements with an higher rank $j > i$, therefore minimal elements of

C are those having the minimum rank. In detail, $.^{\mathcal{I}}$ 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 the $\mathbf{T}$ operator, we have

$$(\mathbf{T}(C))^{\mathcal{I}} = Min_{<}(C^{\mathcal{I}}),$$

where $Min_{<}(C^{\mathcal{I}}) = \{x \in C^{\mathcal{I}} \mid \nexists y \in C^{\mathcal{I}} \text{ s.t. } y < x\}$.

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 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 [Giordano et al. (2015)] 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 atypical instances of a concept*. The resulting logic 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 [Lehmann and Magidor (1992)] for the propositional logic.

The nonmonotonic semantics of $\mathcal{ALC} + \mathbf{T_R}$ 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.

Definition 2.1 (Rank of a domain element $k_{\mathcal{M}}(x)$)**.** 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 $x_0 < \cdots < x$ from x to a minimal x_0 (i.e. such that there is no x' such that $x' < x_0$).

Definition 2.2 (Minimal models)**.** Given $\mathcal{M}_1 = \langle \Delta^{\mathcal{I}_1}, <_1, .^{\mathcal{I}_1} \rangle$ and $\mathcal{M}_2 = \langle \Delta^{\mathcal{I}_2}, <_2, .^{\mathcal{I}_2} \rangle$ we say that $\mathcal{M}_1$ is preferred to $\mathcal{M}_2$ if:

- $\Delta^{\mathcal{I}_1} = \Delta^{\mathcal{I}_2}$
- $C^{\mathcal{I}_1} = C^{\mathcal{I}_2}$ for all concepts C
- for all $x \in \Delta^{\mathcal{I}_1}$, it holds that $k_{\mathcal{M}_1}(x) \leq k_{\mathcal{M}_2}(x)$ whereas there exists $y \in \Delta^{\mathcal{I}_1}$ such that $k_{\mathcal{M}_1}(y) < k_{\mathcal{M}_2}(y)$.

Given a knowledge base, we say that $\mathcal{M}$ is a minimal model if it is a model satisfying the knowledge base and there is no $\mathcal{M}'$ model satisfying it such that $\mathcal{M}'$ is preferred to $\mathcal{M}$.

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[3]. A query F is minimally entailed from a KB if it holds in all minimal canonical models of KB.

Example 2.3. Let us consider and extend the example in the Introduction about athletes and sumo wrestlers. In the logic $\mathcal{ALC} + \mathbf{T_R}$ we can have a knowledge base whose TBox contains the following inclusions:

[3]In Theorem 10 in [Giordano et al. (2015)] the authors have shown that for any consistent KB there exists a finite minimal canonical model of KB.

$SumoWrestler \sqsubseteq Athlete$
$Athlete \sqsubseteq HumanBeing$
$\mathbf{T}(Athlete) \sqsubseteq Fit$
$\mathbf{T}(SumoWrestler) \sqsubseteq \neg Fit$
$\mathbf{T}(Athlete) \sqsubseteq YoungPerson$

where standard inclusions are intended as usual in standard $\mathcal{ALC}$: all sumo wrestlers are athletes, and all athletes are human beings. Typicality inclusions represent that, respectively, usually, athletes are fit, typical sumo wrestlers are not fit, and, normally, athletes are young persons.

The ABox contains the following facts about Roberto and Hiroyuki:

$Athlete(roberto)$
$SumoWrestler(hiroyuki)$

used to represent that Roberto is an athlete, whereas Hiroyuky is a sumo wrestler. Concerning its nonmonotonic reasoning capabilities, the logic $\mathcal{ALC} + \mathbf{T_R}$ allows one to infer from the above knowledge base that:

$\mathbf{T}(Athlete) \sqsubseteq \neg SumoWrestler$
$\mathbf{T}(Athlete \sqcap Bald) \sqsubseteq Fit$ (*irrelevance*)

the last one stating that being bald is irrelevant with respect to being in fit. Furthermore, since we know that Hiroyuki is a sumo wrestler and Roberto is an athlete, we can infer the following facts:

$Fit(roberto)$
$\neg Fit(hiroyuky)$ (*specificity*)

Observe that, in the last one, the logic gives preference to the most specific information (Hiroyuky is both an athlete and a sumo wrestler).

In [Giordano et al. (2015)] it is shown that query entailment in $\mathcal{ALC} + \mathbf{T_R}$ is in EXPTIME.

2.3. Probabilistic DLs: the DISPONTE semantics

A probabilistic extension of Description Logics under the distribution semantics is proposed in [Riguzzi et al. (2015a)]. In this approach, called DISPONTE, the authors propose the integration of probabilistic information with DLs based on the distribution semantics for probabilistic logic programs [Sato (1995)]. 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 (intended as *degrees of belief*), assuming that each axiom is independent from each other. 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 a query is obtained by marginalizing the joint distribution of the query and the worlds.

As an example, consider the following variant of the knowledge base inspired by the people and pets ontology in [Riguzzi et al. (2015a)]:

$0.3 :: \exists hasAnimal.Pet \sqsubseteq NatureLover$ (1)
$0.6 :: Cat \sqsubseteq Pet$ (2)
$0.9 :: Cat(tom)$ (3)
$hasAnimal(kevin, tom)$ (4)

The inclusion (1) expresses that individuals that own a pet are nature lovers with a 30%

probability, whereas (2) is used to state that cats are pets with probability 60%. The ABox fact (3) represents that Tom is a cat with probability 90%. Inclusions (1), (2) and (3) are *probabilistic* axioms, whereas (4) is a *certain* axiom, that must always hold. The KB has the following eight possible worlds:

$$\{((1),0),((2),0),((3),0)\}$$
$$\{((1),0),((2),0),((3),1)\}$$
$$\{((1),0),((2),1),((3),0)\}$$
$$\{((1),0),((2),1),((3),1)\}$$
$$\{((1),1),((2),0),((3),0)\}$$
$$\{((1),1),((2),0),((3),1)\}$$
$$\{((1),1),((2),1),((3),0)\}$$
$$\{((1),1),((2),1),((3),1)\}$$

representing all possible combinations of considering/not considering each probabilistic axiom. For instance, the world $\{((1),1),((2),0),((3),1)\}$ represents the situation in which we have that (1) and (3) hold, i.e. $\exists hasAnimal.Pet \sqsubseteq NatureLover$ and $Cat(tom)$, whereas (2) does not. The query

$$NatureLover(kevin)$$

is true only in the last world, i.e. having that (1), (2) and (3) are all true, whereas it is false in all the other ones. The probability of such a query is $P(NatureLover(kevin)) = 0.3 \times 0.6 \times 0.9 = 0.162$.

It is worth noticing that, in $\mathbf{T}^{\mathsf{CL}}$, we do not employ the whole characteristics of the DISPONTE semantics. In particular, as we will describe in detail in the next section, as far as the inferential capabilities are concerned, we exclusively adopt $\mathcal{ALC} + \mathbf{T_R}$, whereas we use DISPONTE only as a (necessary) ingredient to generate all different knowledge bases obtained by considering different subsets of typicality inclusions.

3. $\mathbf{T}^{\mathsf{CL}}$: A Logic for Commonsense Concept Combination

In this section, we introduce a new nonmonotonic Description Logic $\mathbf{T}^{\mathsf{CL}}$ that combines the semantics based on the rational closure of $\mathcal{ALC} + \mathbf{T_R}$ [Giordano et al. (2015)] with the DISPONTE semantics [Riguzzi et al. (2015a,b)] of probabilistic DLs.

By taking inspiration from [Lieto et al. (2015, 2017)], in our representational assumptions we consider two different types of properties associated to a given concept: rigid and typical. Rigid properties are those defining a concept, e.g. $C \sqsubseteq D$ (all Cs are Ds). Typical properties are represented by inclusions equipped by a degree of belief expressed through probabilities like in the DISPONTE Semantics. Additionally, as mentioned, we employ insights coming from the cognitive science for the determination of a dominance effect between the concepts to be combined, distinguishing between concept HEAD and MODIFIER. Since the conceptual combination is usually expressed via natural language we consider the following common situations: in a combination ADJECTIVE - NOUN (for instance, *red apple*) the HEAD is represented by the NOUN (*apple*) and the modifier by the ADJECTIVE (*red*). In the more complex case of NOUN-NOUN combinations (for instance, *pet fish*) usually the HEAD is represented by the last expressed concept (*fish* in this case). As we will see, however, in the NOUN-NOUN case (i.e. the one we will take into account in this paper) does not exists a clear rule to follow [4].

[4]It is worth-noting that a general framework for the automatic identification of a HEAD/MODIFIER combination is currently

The language of $\mathbf{T}^{\mathsf{CL}}$ extends the basic DL $\mathcal{ALC}$ by *typicality inclusions* of the form $\mathbf{T}(C) \sqsubseteq D$ equipped by a real number $p \in (0.5, 1]$ – observe that the extreme 0.5 is not included – representing its degree of belief, whose meaning is that "we believe with degree/probability p that, normally, Cs are also D" [5].

Definition 3.1 (Language of $\mathbf{T}^{\mathsf{CL}}$)**.** We consider an alphabet of concept names $\mathtt{C}$, of role names $\mathtt{R}$, and of individual constants $\mathtt{O}$. Given $A \in \mathtt{C}$ and $R \in \mathtt{R}$, we define:

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

We define a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ where:
• $\mathcal{R}$ is a finite set of rigid properties of the form $C \sqsubseteq D$;
• $\mathcal{T}$ is a finite set of typicality properties of the form

$$p \ :: \ \mathbf{T}(C) \sqsubseteq D$$

where $p \in (0.5, 1] \subseteq \mathbb{R}$ is the degree of belief of the typicality inclusion;
• $\mathcal{A}$ is the ABox, i.e. a finite set of formulas of the form either $C(a)$ or $R(a, b)$, where $a, b \in \mathtt{O}$ and $R \in \mathtt{R}$.

Example 3.2. Let us consider and extend the previous example about athletes and sumo wrestlers. In the logic $\mathbf{T}^{\mathsf{CL}}$ we can have a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ as follows:

$\mathcal{R}$:

- $SumoWrestler \sqsubseteq Athlete$
- $Athlete \sqsubseteq HumanBeing$

$\mathcal{T}$:

- $0.8 \ :: \ \mathbf{T}(Athlete) \sqsubseteq Fit$
- $0.8 \ :: \ \mathbf{T}(SumoWrestler) \sqsubseteq \neg Fit$
- $0.95 \ :: \ \mathbf{T}(Athlete) \sqsubseteq YoungPerson$

$\mathcal{A}$:

- $Athlete(roberto)$
- $SumoWrestler(hiroyuki)$

Rigid properties of $\mathcal{R}$ are intended as usual in standard $\mathcal{ALC}$: all sumo wrestlers are athletes, and all athletes are human beings. Typicality properties of $\mathcal{T}$ represent the following facts, respectively:

- usually, athletes are fit, and this fact has a degree of belief of 80%;
- typical sumo wrestlers are not fit with a degree of belief of 80%;
- we believe with a degree of 95% in the fact that, normally, athletes are young persons.

The ABox facts in $\mathcal{A}$ are used to represent that Roberto is an athlete, whereas Hiroyuky is a sumo wrestler.

We remind that, since we exploit the logic of typicality $\mathcal{ALC} + \mathbf{T}_{\mathbf{R}}$, our logic $\mathbf{T}^{\mathsf{CL}}$ inherits its nonmonotonic reasoning capabilities.

As mentioned above, we do not avoid typicality inclusions with degree 1. Indeed, an inclusion $1 \ :: \ \mathbf{T}(C) \sqsubseteq D$ means that there is no uncertainty about a given typicality inclusion.

---

not available in literature. In this work we will take for granted that some methods for the correct identification of these pairs exist and we will focus on the reasoning part.

[5] The reason why we only allow typicality inclusions equipped with probabilities $p > 0.5$ is detailed in the following.

On the other hand, since the very cognitive notion of typicality [Rosch (1975)] derives from that one of probability distribution [6], and this latter notion is also intrinsically connected to the one concerning the level of uncertainty/degree of belief associated to typicality inclusions (i.e. typical knowledge is known to come with a low degree of uncertainty [Lawry and Tang (2009)]), we only allow typicality inclusions equipped with degrees of belief $p > 0.5$[7]. Therefore, in our effort of integrating two different semantics – typicality based logic and DISPONTE – the choice of having degrees of belief/probabilities higher than 0.5 for typicality inclusions seems to be the only one compliant with both the formalisms. In fact, despite the DISPONTE semantics [Riguzzi et al. (2015b)] allows to assign also low degrees of belief to standard inclusions, in the logic $\mathbf{T}^{\mathsf{CL}}$, for what explained above, it would be at least counter-intuitive to also allow low degrees of belief for typicality inclusions (simply because typicality inclusions with high uncertainty do not describe any typical knowledge). For example, the logic $\mathbf{T}^{\mathsf{CL}}$ does not allow an inclusion like $0.3 :: \mathbf{T}(Student) \sqsubseteq YoungPerson$, that could be interpreted as "normally, students are not young people". Please, note that this is not a limitation of the expressivity of the logic $\mathbf{T}^{\mathsf{CL}}$: we can in fact represent properties not holding for typical members of a category, for instance if one needs to represent that typical students are not married, we can have that $0.8 :: \mathbf{T}(Student) \sqsubseteq \neg Married$, rather than $0.2 :: \mathbf{T}(Student) \sqsubseteq Married$.

Following the DISPONTE semantics, in $\mathbf{T}^{\mathsf{CL}}$ each typicality inclusion is independent from each other. This avoids the problem of dealing with degrees of inconsistent inclusions. Let us consider the following knowledge base:

$WorkingStudent \sqsubseteq Student$
$(i)\ 0.8 :: Student \sqsubseteq \neg WorkingTaxPayer$
$(ii)\ 0.9 :: WorkingStudent \sqsubseteq WorkingTaxPayer$

Also in the scenarios where both the conflicting typical inclusions (i) and (ii) are considered, the two degrees describe, respectively, that we believe that typical students do not pay working taxes with a degree of 80%, and that working students normally pay working taxes with a degree of 90%, then those inclusions are both acceptable due to the independence assumption. The two degrees will contribute to a definition of probability of such scenario (as we will describe in Definition 3.8). It is worth noticing that the underlying logic of typicality allows us to get for free the correct way of reasoning in this case, namely if the ABox contains the information that $Mark$ is a working student, we obtain that he pays working taxes, i.e. $WorkingTaxPayer(Mark)$.

A model $\mathcal{M}$ in the logic $\mathbf{T}^{\mathsf{CL}}$ extends standard $\mathcal{ALC}$ models by a preference relation among domain elements as in the logic of typicality [Giordano et al. (2015)]. 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[8]. An element $x \in \Delta^{\mathcal{I}}$ is a *typical instance* of some concept C if $x \in C^{\mathcal{I}}$ and there is no C-element in $\Delta^{\mathcal{I}}$ *more normal* than x. Formally:

[6] In the sense that prototypes are usually intended as statistically representative members of a category (e.g. the prototype of *Bird* in, let's say, Europe is different with respect to the prototype of *Bird* in New Zealand, i.e. a kiwi which do not fly, since the types of birds encountered by these two populations are statistically different and therefore the learned typical members differ as well). This assumption is also reflected in computational frameworks of cognitive semantics where prototypes for conceptual representations are naturally calculated/learned as centroids in vector space models of meaning[Gärdenfors (2004, 2014)].

[7] It is worth noticing that in our logic, the uncertain/graded component of typicality is captured by the *ranked* semantics underlying the operator $\mathbf{T}$ in the logic $\mathcal{ALC}+\mathbf{T_R}$. On the other hand, the epistemic uncertainty is modelled by the interpretation of probabilities of the DISPONTE semantics.

[8] It could be possible to consider an alternative semantics whose models are equipped with multiple preference relations, whence with multiple typicality operators. In this case, it should be possible to distinguish different aspects of exceptionality, however the approach based on a single preference relation in [Giordano et al. (2015)] ensures good computational properties (reasoning in the resulting nonmonotonic logic $\mathcal{ALC} + \mathbf{T_R}$ has the same complexity of the standard $\mathcal{ALC}$), whereas adopting multiple preference relations could lead to higher complexities.

Definition 3.3 (Model of $\mathbf{T}^{\mathsf{CL}}$)**.** A 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 atomic concept C to $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and each role R to $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and is extended to complex concepts as follows:
  - $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
  - $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$
  - $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$
  - $(\exists R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \exists (x, y) \in R^{\mathcal{I}} \text{ such that } y \in C^{\mathcal{I}}\}$
  - $(\forall R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \forall (x, y) \in R^{\mathcal{I}} \text{ we have } y \in C^{\mathcal{I}}\}$
  - $(\mathbf{T}(C))^{\mathcal{I}} = Min_{<}(C^{\mathcal{I}})$, where $Min_{<}(C^{\mathcal{I}}) = \{x \in C^{\mathcal{I}} \mid \nexists y \in C^{\mathcal{I}} \text{ s.t. } y < x\}$.

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 [Giordano et al. (2015)]: the rank of x is the length of the longest chain $x_0 < \cdots < x$ from x to a minimal x_0, i.e. such that there is no x' such that $x' < x_0$. The rank function $k_{\mathcal{M}}$ and $<$ can be defined from each other by letting $x < y$ if and only if $k_{\mathcal{M}}(x) < k_{\mathcal{M}}(y)$.

Definition 3.4 (Model satisfying a knowledge base in $\mathbf{T}^{\mathsf{CL}}$)**.** Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB. Given a model $\mathcal{M} = \langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$, we assume that $.^{\mathcal{I}}$ is extended to assign a domain element $a^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ to each individual constant a of O. We say that:

- $\mathcal{M}$ satisfies $\mathcal{R}$ if, for all $C \sqsubseteq D \in \mathcal{R}$, we have $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$;
- $\mathcal{M}$ satisfies $\mathcal{T}$ if, for all $q \ :: \ \mathbf{T}(C) \sqsubseteq D \in \mathcal{T}$, we have that[9] $\mathbf{T}(C)^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, i.e. $Min_{<}(C^{\mathcal{I}}) \subseteq D^{\mathcal{I}}$;
- $\mathcal{M}$ satisfies $\mathcal{A}$ if, for all assertion $F \in \mathcal{A}$, if $F = C(a)$ then $a^{\mathcal{I}} \in C^{\mathcal{I}}$, otherwise if $F = R(a, b)$ then $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$.

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 resulting logic is monotonic. As already mentioned in Section 2, in order to perform useful nonmonotonic inferences, in [Giordano et al. (2015)] 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 atypical instances of a concept*. The resulting logic 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 [Lehmann and Magidor (1992)] for the propositional logic. This nonmonotonic semantics 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

[9] It is worth noticing that here the degree q does not play any role. Indeed, a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ holds in a model only if it satisfies the semantic condition of the underlying DL of typicality, i.e. minimal (typical) elements of C are elements of D. The degree of belief q will have a crucial role in the application of the distributed semantics, allowing the definition of scenarios as well as the computation of their probabilities.

the individuals that enjoy properties that are consistent with KB. This is needed when reasoning about the rank of the concepts: it is important to have them all represented. A query F is minimally entailed from a KB if it holds in all minimal canonical models of KB. In [Giordano et al. (2015)] it is shown that query entailment in the nonmonotonic $\mathcal{ALC} + \mathbf{T_R}$ is in EXPTIME.

Definition 3.5 (Entailment)**.** Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB and let F be either $C \sqsubseteq D$ (C could be $\mathbf{T}(C')$) or $C(a)$ or $R(a, b)$. We say that F follows from $\mathcal{K}$ if, for all minimal $\mathcal{M}$ satisfying $\mathcal{K}$, then $\mathcal{M}$ also satisfies F.

Let us now define the notion of *scenario* of the composition of concepts. Intuitively, a scenario is a knowledge base obtained by adding to all rigid properties in $\mathcal{R}$ and to all ABox facts in $\mathcal{A}$ only *some* typicality properties. More in detail, we define an *atomic choice* on each typicality inclusion, then we define a *selection* as a set of atomic choices in order to select which typicality inclusions have to be considered in a scenario.

Definition 3.6 (Atomic choice)**.** Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 \ :: \mathbf{T}(C_1) \sqsubseteq D_1, \ldots, E_n = q_n \ :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ we define (E_i, k_i) an *atomic choice*, where $k_i \in \{0, 1\}$.

Definition 3.7 (Selection)**.** Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 \ :: \mathbf{T}(C_1) \sqsubseteq D_1, \ldots, E_n = q_n \ :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ and a set of atomic choices ν, we say that ν is a *selection* if, for each E_i, one decision is taken, i.e. either $(E_i, 0) \in \nu$ and $(E_i, 1) \notin \nu$ or $(E_i, 1) \in \nu$ and $(E_i, 0) \notin \nu$ for $i = 1, 2, \ldots, n$. The probability of ν is $P(\nu) = \prod_{(E_i,1)\in\nu} q_i \prod_{(E_i,0)\in\nu} (1 - q_i)$.

Definition 3.8 (Scenario)**.** Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 \ :: \mathbf{T}(C_1) \sqsubseteq D_1, \ldots, E_n = q_n \ :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ and given a selection σ, we define a *scenario* $w_\sigma = \langle \mathcal{R}, \{E_i \mid (E_i, 1) \in \sigma\}, \mathcal{A} \rangle$. We also define the probability of a scenario w_σ as the probability of the corresponding selection, i.e. $P(w_\sigma) = P(\sigma)$. Last, we say that a scenario is *consistent* with respect to $\mathcal{K}$ when it admits a model in the logic $\mathbf{T}^{\mathsf{CL}}$ satisfying $\mathcal{K}$.

We denote with $\mathcal{W}_\mathcal{K}$ the set of all scenarios, essentially the set of knowledge bases obtained by considering all possible subsets of typicality inclusions in $\mathcal{T}$. It immediately follows that the probability of a scenario $P(w_\sigma)$ is a probability distribution over scenarios, that is to say $\sum_{w\in\mathcal{W}_\mathcal{K}} P(w) = 1$.

Given a KB $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ and given two concepts C_H and C_M occurring in $\mathcal{K}$, our logic allows defining the compound concept C as the combination of the HEAD C_H and the MODIFIER C_M, where the typical properties of the form $\mathbf{T}(C) \sqsubseteq D$ (or, equivalently, $\mathbf{T}(C_H \sqcap C_M) \sqsubseteq D$) to ascribe to the concept C are obtained in the set of scenarios that:

(1) are consistent with respect to $\mathcal{K}$ in presence of at least a C-element, in other words the knowledge base extending $\mathcal{K}$ with the properties ascribed to the combined concept C in the scenario, i.e. $\langle \mathcal{R}, \mathcal{T} \cup \{\mathbf{T}(C_H \sqcap C_M) \sqsubseteq D \mid \text{either } \mathbf{T}(C_H) \sqsubseteq D \in w_\sigma \text{ or } \mathbf{T}(C_M) \sqsubseteq D \in w_\sigma\}, \mathcal{A} \cup \{(C_H \sqcap C_M)(x)\} \rangle$, where x does not occur in $\mathcal{A}$, admits a model in $\mathbf{T}^{\mathsf{CL}}$;

(2) are not trivial, i.e. the scenarios with the highest probability considering either *all* properties that can be consistently ascribed to C are discarded or *all* properties of the HEAD that can be consistently ascribed to C are discarded [10];

[10] This choice is motivated by the challenges provided by the task of commonsense conceptual combination itself: in order to generate plausible novel compounds it is necessary to maintain a certain level of "surprise" in the combination, since obvious inheritance of attributes does not have any explanatory power for human-like and human-level concept combination [Hampton (1987)]. For this reason, both scenarios inheriting all the properties of the two concepts and all the properties of the HEAD are

(3) are those giving preference to the typical properties of the HEAD C_H (with respect to those of the MODIFIER C_M) with the highest probability, that is to say a scenario w is discarded if, in case of conflicting properties D and $\neg D$, w contains an inclusion $p_1 :: \mathbf{T}(C_M) \sqsubseteq \neg D$ whereas it does not include another inclusion $p_2 :: \mathbf{T}(C_H) \sqsubseteq D$.

In order to select the resulting scenarios we apply points 1, 2, and 3 above to blocks of scenarios with the same probability, in decreasing order starting from the highest one. More in detail, we first discard all the inconsistent scenarios, then we consider the remaining (consistent) ones in decreasing order by their probabilities. We then consider the blocks of scenarios with the same probability, and we proceed as follows:

- we discard those considered as *trivial*, consistently inheriting all the properties from the HEAD (therefore, also scenarios inheriting all the properties of HEAD and MODIFIER are discarded) from the starting concepts to be combined;
- among the remaining ones, we discard those inheriting properties from the MODIFIER in conflict with properties that could be consistently inherited from the HEAD;
- if the set of scenarios of the current block is empty, i.e. all the scenarios have been discarded either because trivial or because preferring the MODIFIER, we repeat the procedure by considering the block of scenarios, all having the immediately lower probability;
- the set of remaining scenarios are those selected by the logic $\mathbf{T}^{\mathsf{CL}}$.

More formally, our mechanism is described in Algorithm 1. Please note that this block-based procedure extends a previously developed method that simply selected the consistent scenarios with the probability range immediately lower to the non-trivial ones [Lieto and Pozzato (2018)]. Notice also that, in the initial knowledge base $\mathcal{K}$, we have that the set of typicality inclusions is $\mathcal{T}' \cup \{q_1 :: \mathbf{T}(C_H) \sqsubseteq D_1, \ldots, q_k :: \mathbf{T}(C_H) \sqsubseteq D_k, q_{k+1} :: \mathbf{T}(C_M) \sqsubseteq D_{k+1}, q_n :: \mathbf{T}(C_M) \sqsubseteq D_n\}$, where $\mathcal{T}'$ does not contain neither inclusions of the HEAD of the form $q_i :: \mathbf{T}(C_H) \sqsubseteq D_i$ nor inclusions of the MODIFIER of the form $q_i :: \mathbf{T}(C_M) \sqsubseteq D_i$ (Algorithm 1, line 2).

Lastly, we define the ultimate output of our mechanism: a knowledge base in the logic $\mathbf{T}^{\mathsf{CL}}$ whose set of typicality properties is enriched by those of the compound concept C. Given a scenario w satisfying the above properties, we define the properties of C as the set of inclusions $p :: \mathbf{T}(C) \sqsubseteq D$, for all $\mathbf{T}(C) \sqsubseteq D$ that are entailed (Definition 3.5) from w in the logic $\mathbf{T}^{\mathsf{CL}}$. The probability p is such that:

- if $\mathbf{T}(C_H) \sqsubseteq D$ is entailed from w, that is to say D is a property inherited either from the HEAD (or from both the HEAD and the MODIFIER), then p corresponds to the probability of such inclusion of the HEAD in the initial knowledge base, i.e. $p : \mathbf{T}(C_H) \sqsubseteq D \in \mathcal{T}$;
- otherwise, i.e. $\mathbf{T}(C_M) \sqsubseteq D$ is entailed from w, then p corresponds to the probability of such inclusion of a MODIFIER in the initial knowledge base, i.e. $p : \mathbf{T}(C_M) \sqsubseteq D \in \mathcal{T}$.

The knowledge base obtained as the result of combining concepts C_H and C_M into the compound concept C is called *C-revised* knowledge base, and it is defined as follows:

$$\mathcal{K}_C = \langle \mathcal{R}, \mathcal{T} \cup \{p : \mathbf{T}(C) \sqsubseteq D\}, \mathcal{A} \rangle,$$

for all D such that either $\mathbf{T}(C_H) \sqsubseteq D$ is entailed in w or $\mathbf{T}(C_M) \sqsubseteq D$ is entailed in w by

discarded. In this respect, the typicality-based inheritance procedure of our logic falls within the so-called *functional compositionality*, as introduced in [Pelletier (2017)].

Algorithm 1 Algorithm for concept combination

procedure CONCEPTCOMBINATION($\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle, C_H, C_M$)
  let $\mathcal{T}$ be $\mathcal{T}' \cup \{q_1 :: \mathbf{T}(C_H) \sqsubseteq D_1, \ldots, q_k :: \mathbf{T}(C_H) \sqsubseteq D_k, q_{k+1} :: \mathbf{T}(C_M) \sqsubseteq D_{k+1}, q_n :: \mathbf{T}(C_M) \sqsubseteq D_n\}$
  Sel $\leftarrow$ compute all the 2^n possible selections
  Snr $\leftarrow \emptyset$ ▷ Compute all 2^n scenarios
  for each $\nu \in$ Sel **do**
    $P(\nu) \leftarrow 1$ ▷ Compute the probability of the current selection/scenario
    for $i = 1, \ldots, n$ **do**
      if $\nu_i = 1$ **then**
        $P(\nu) \leftarrow P(\nu) \times q_i$
      else
        $P(\nu) \leftarrow P(\nu) \times (1 - q_i)$
    $\mathcal{T}_\nu \leftarrow \emptyset$ ▷ Build the scenario corresponding to ν
    for $i = 1, \ldots, n$ **do**
      if $\nu_i = 1$ **then**
        $\mathcal{T}_\nu \leftarrow \mathcal{T}_\nu \cup \{q_i :: \mathbf{T}(C_H \sqcap C_M) \sqsubseteq D_i\}$
    Snr $\leftarrow$ Snr $\cup \{\langle \mathcal{R}, \mathcal{T}' \cup \mathcal{T}_\nu, \mathcal{A} \rangle\}$
  ConsSnr $\leftarrow \emptyset$ ▷ Discard inconsistent scenarios (reasoning in $\mathcal{ALC} + \mathbf{T_R}$)
  for each $\langle \mathcal{R}, \mathcal{T} \cup \mathcal{T}_\nu, \mathcal{A} \rangle \in$ Snr **do**
    if $\langle \mathcal{R}, \mathcal{T} \cup \mathcal{T}_\nu, \mathcal{A} \cup \{(C_H \sqcap C_M)(x)\} \rangle$ is consistent in $\mathcal{ALC} + \mathbf{T_R}$ **then** ▷ x does not occur in $\mathcal{A}$
      ConsSnr $\leftarrow$ ConsSnr $\cup \{\langle \mathcal{R}, \mathcal{T} \cup \mathcal{T}_\nu, \mathcal{A} \rangle\}$
  Ord $\leftarrow$ order scenarios in ConsSnr by probabilities $P(\nu)$ in a decreasing order
  ResultingSnr $\leftarrow \emptyset$ ▷ The set ResultingSnr will contain the selected scenarios
  while ResultingSnr $== \emptyset$ **do** ▷ Continue with the next block of scenarios
    $w \leftarrow$ first scenario in Ord
    Max $\leftarrow P(w)$ ▷ Highest probability in Ord
    CurrentBlock $\leftarrow \{w\}$ ▷ Build the current set scenarios with the highest probability
    while $P(w) ==$ Max **do**
      $w \leftarrow$ next scenario in Ord ▷ w is not removed from Ord
      if $P(w) ==$ Max **then**
        CurrentBlock $\leftarrow$ CurrentBlock $\cup \{w\}$
        remove w from Ord
    for each $w \in$ CurrentBlock **do**
      if w not contains all properties of C_H **then** ▷ Trivial scenario to be discarded
        if CONFLICTHEADMODIFIER(w,$\mathcal{K}$,C_H,C_M) $==$ *false* **then**
          ▷ Scenario preferring MODIFIER to HEAD (see Algorithm 2). Scenario to be discarded
          ResultingSnr $\leftarrow$ ResultingSnr $\cup \{w\}$ ▷ Selected scenario
  return ResultingSnr

Definition 3.5, and p is defined as above.

Let us now define the probability that a query is entailed from a C-revised knowledge base. We restrict our concern to ABox facts. The intuitive idea is that, given a query F of the form $A(a)$ and its associated probability p, the probability of F is the product of p and the

Algorithm 2 Algorithm checking for a scenario giving preference to the HEAD

```
1: procedure ConflictHeadModifier(w,K,C_H,C_M)
2:     for each q_j :: T(C_M) ⊑ D ∈ w do
3:         for each q_k :: T(C_H) ⊑ E ∈ K such that q_k :: T(C_H) ⊑ E ∉ w do
4:             if K ∪ {D(x), E(x)} is inconsistent in ALC + T_R, x ∉ K then
5:                                   ▷ Properties D and E are contradictory in ALC + T_R
6:                 return true
7:     return false
```

probability of the inclusion in the C-revised knowledge base which is responsible for that.

Definition 3.9 (Probability of query entailment)**.** Given a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, the C-revised knowledge base $\mathcal{K}_C$, a query $A(a)$ and its probability $p \in (0, 1]$, we define the probability of the entailment of the query $A(a)$ from $\mathcal{K}_C$, denoted as $\mathbb{P}(A(a), p)$ as follows:

- $\mathbb{P}(A(a), p) = 0$, if $A(a)$ is *not* entailed from $\mathcal{K}_C$;
- $\mathbb{P}(A(a), p) = p \times q$, where either $q :: \mathbf{T}(C) \sqsubseteq A$ belongs to $\mathcal{K}_C$ or $q :: \mathbf{T}(C) \sqsubseteq D$ belongs to $\mathcal{K}_C$ and $D \sqsubseteq A$ is entailed from $\mathcal{R}$ in standard $\mathcal{ALC}$, otherwise.

We conclude this section by showing that reasoning in $\mathbf{T}^{\mathsf{CL}}$ remains in the same complexity class of standard $\mathcal{ALC}$ Description Logics.

Theorem 3.10. *Reasoning in* $\mathbf{T}^{\mathsf{CL}}$ *is* EXPTIME*-complete.*

Proof. For the completeness, let n be the size of KB, then the number of typicality inclusions is $O(n)$. It is straightforward to observe that we have an exponential number of different scenarios, for each one we need to check whether the resulting KB is consistent in $\mathcal{ALC} + \mathbf{T_R}$ which is EXPTIME-complete. Hardness immediately follows form the fact that $\mathbf{T}^{\mathsf{CL}}$ extends standard $\mathcal{ALC}$. Reasoning in the revised knowledge base relies on reasoning in $\mathcal{ALC} + \mathbf{T_R}$, therefore we can conclude that reasoning in $\mathbf{T}^{\mathsf{CL}}$ is EXPTIME-complete. □

4. Applications of the logic $\mathbf{T}^{\mathsf{CL}}$

We propose three different types of examples adopting the logic $\mathbf{T}^{\mathsf{CL}}$, along with its embedded HEAD-MODIFIER heuristic, to model the phenomenon of typicality-based conceptual combination. In the first case (*pet fish*) we show how our logic is able to handle, under certain assumptions, this concept composition which is problematic for other formalisms. In the second case (*Linda the feminist bank teller*) we show how $\mathbf{T}^{\mathsf{CL}}$ is able to model the well known conjunction fallacy problem [Tversky and Kahneman (1983)]. In the third case (*stone lion*) we show how our logic is also able to account for complex form of metaphorical concept combination. All these examples do not come *ex-abrupto*, since they represent classical challenging cases to model in the field of cognitive science and cognitive semantics (see e.g. [Lewis and Lawry (2016)]) and have been showed in that past problematic to model by adopting other kinds of logics (for example fuzzy logic, [Hampton (2011); Osherson and Smith (1981); Smith and Osherson (1984)]).

In addition, we exploit $\mathbf{T}^{\mathsf{CL}}$ to present an example of a possible application in the area of creative generation of new characters. Finally, we show that the logic $\mathbf{T}^{\mathsf{CL}}$ can be iteratively applied to combine concepts already resulting from the combination of concepts. This type of iterative process has been never provided in previous formalizations trying to address similar

or the very same phenomena, (e.g. in [Eppe et al. (2018); Lewis and Lawry (2016)]). We show that the procedures provided in $\mathbf{T}^{\mathsf{CL}}$ are robust and consistent enough also for dealing with higher, iterative, levels of prototype-based compositionality.

4.1. Pet Fish

In this section we exploit, in two different setups, the logic $\mathbf{T}^{\mathsf{CL}}$ in order to define the typical properties of the concept *pet fish*, obtained as the combination of the concepts *Pet* and *Fish*. As mentioned before, this represents a well known and paradigmatic example in cognitive science. The problem of combining the prototype of a pet with those of a fish is the following: a typical pet is affectionate and warm, whereas a pet fish is not; on the other hand, as a difference with a typical fish, a pet fish is not greyish, but it inherits its being scaly.

Let us consider, as a first setup, the following situation: let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB, where the ABox $\mathcal{A}$ is empty, the set of rigid inclusions is

$$\mathcal{R} = \{Fish \sqsubseteq \forall livesIn.Water\}$$

and the set of typicality properties $\mathcal{T}$ is as follows:

(1) 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Affectionate$
(2) 0.6 :: $\mathbf{T}(Fish) \sqsubseteq Greyish$
(3) 0.9 :: $\mathbf{T}(Fish) \sqsubseteq Scaly$
(4) 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Warm$
(5) 0.9 :: $\mathbf{T}(Pet) \sqsubseteq \forall livesIn.(\neg Water)$
(6) 0.8 :: $\mathbf{T}(Pet) \sqsubseteq Affectionate$
(7) 0.8 :: $\mathbf{T}(Pet) \sqsubseteq Warm$

By the properties of the typicality operator $\mathbf{T}$, reasoning relying on the underlying $\mathcal{ALC}+\mathbf{T_R}$ we have that

$$(*)\ \mathbf{T}(Pet \sqcap Fish) \sqsubseteq \forall livesIn.Water.$$

Indeed, $Fish \sqsubseteq \forall livesIn.Water$ is a rigid property, which is always preferred to a typical one: in this case, additionally, the rigid property is also associated to the HEAD element fish. Therefore, this element is reinforced.

Since $\mid \mathcal{T} \mid = 7$, we have $2^7 = 128$ different scenarios. We can observe that some of them are not consistent[11], more precisely those

(i) containing the inclusion 5, thus contradicting $(*)$;
(ii) containing both inclusions 1 and 6;
(iii) containing both inclusions 4 and 7.

It is worth noticing that this setting of the example represents one of worst cases in our analysis: indeed, the probabilities associated to the properties in $\mathcal{T}$ related to the MODIFIER are not lower than the ones associated to the properties in $\mathcal{T}$ related to the HEAD[12].

The scenario with the highest probability (up to 20%) is both trivial and inconsistent: indeed, since probabilities p_i equipping typicality inclusions are such that $p_i > 0.5$ by defini-

[11] The inconsistency arises when the knowledge base is extended by two contradicting properties for the combined concept, for instance the knowledge base extended by both $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq Warm$ and $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq \neg Warm$ would be consistent only if there are no pet fishes.

[12] It is worth noticing that the logic $\mathbf{T}^{\mathsf{CL}}$ would select the same scenario also in the more challenging situation in which degrees of properties of the HEAD are strictly lower, for instance in case inclusion 1 would be replaced in $\mathcal{T}$ by 0.7 :: $\mathbf{T}(Fish) \sqsubseteq \neg Affectionate$.

tion, we immediately have that the higher is the number of inclusions belonging to a scenario the higher is the associated probability. Since typicality inclusions introduce properties that are pairwise inconsistent, it follows that such scenarios must be discarded (Algorithm 1, from line 17).

As described in the previous section, we consider the other blocks of consistent scenarios considering their probabilities in descending order (Algorithm 1, lines from 21). Figure 1 shows the $2^7 = 128$ different scenarios, one row for each scenario.

All scenarios with probabilities ranging from 19.907% down to 0.164% are inconsistent. The first valid block contains scenarios whose probability is 0.138%: the four consistent scenarios of this block, however, are discarded. Indeed, they either contain the inclusions 6 but not 1 or 7 but not 4, namely they give preference to the MODIFIER concerning a conflicting property of the HEAD, or they are trivial, i.e. they inherit all the properties of the HEAD (Algorithm 1, lines 32-35).

The next block contains four scenarios with probability of 0.092%. The first two scenarios, again, either contain inclusions 6 and not 1 or contain 7 and not 4, namely again it privileges the MODIFIER with respect to the corresponding negation in the HEAD. Therefore, these scenarios are discarded. The same for the last one, where both 6 and 7 are included rather than 1 and 4. The remaining scenario of this block includes three out of four properties of the HEAD, therefore it is not trivial and it is selected by the logic $\mathbf{T}^{\mathsf{CL}}$ for the composition of the two initial prototypes.

In conclusion, in our proposal, the not trivial scenario defining prototypical properties of a pet fish is defined from the selection $\sigma = \{(1,1),(2,0),(3,1),(4,1),(5,0),(6,0),(7,0)\}$, and contains inclusions 1, 3, and 4, and the resulting scenario w_σ is as follows:

1. 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Affectionate$
3. 0.9 :: $\mathbf{T}(Fish) \sqsubseteq Scaly$
4. 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Warm$

The resulting $Pet \sqcap Fish$-revised knowledge base that the logic $\mathbf{T}^{\mathsf{CL}}$ suggests is as follows:

$$\mathcal{K}_{Pet \sqcap Fish} = \langle \{Fish \sqsubseteq \forall livesIn.Water\}, \mathcal{T}', \emptyset \rangle,$$

where $\mathcal{T}$ is:

- 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Affectionate$
- 0.6 :: $\mathbf{T}(Fish) \sqsubseteq Greyish$
- 0.9 :: $\mathbf{T}(Fish) \sqsubseteq Scaly$
- 0.8 :: $\mathbf{T}(Fish) \sqsubseteq \neg Warm$
- 0.9 :: $\mathbf{T}(Pet) \sqsubseteq \forall livesIn.(\neg Water)$
- 0.8 :: $\mathbf{T}(Pet) \sqsubseteq Affectionate$
- 0.8 :: $\mathbf{T}(Pet) \sqsubseteq Warm$
- 0.8 :: $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq \neg Affectionate$
- 0.9 :: $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq Scaly$
- 0.8 :: $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq \neg Warm$

Notice that in our logic $\mathbf{T}^{\mathsf{CL}}$, adding a new inclusion $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq Red$, would not be problematic. (i.e. this means that our formalism is able to tackle the phenomenon of prototypical *attributes emergence* for the new compound concept, a well established effect within the cognitive science literature [Hampton (1987)]).

Let us now consider the PET FISH problem by adopting a starting knowledge base whose inclusions are equipped with different probabilities. In particular: let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a the

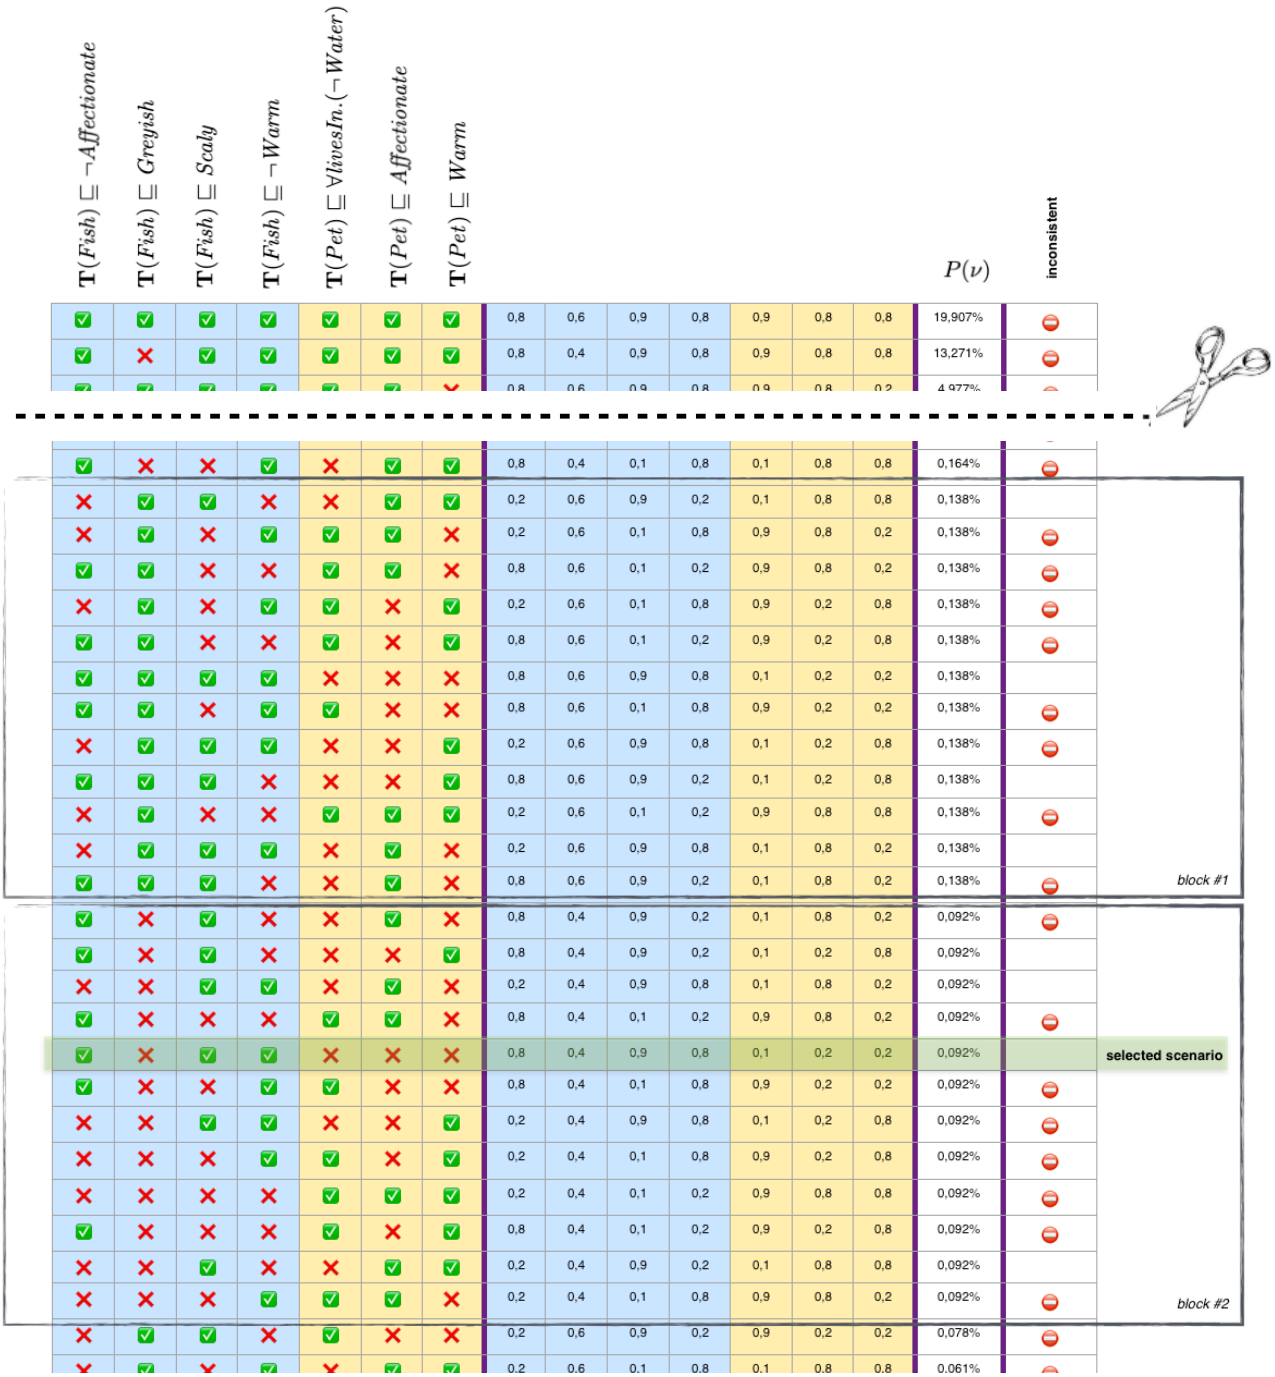

Figure 1.: Scenarios of the combination of Pet and $Fish$. For the sake of readability we cut off the first part of the figure since the first blocks of the combination, i.e. 52 scenarios, are all inconsistent.

KB, where the ABox $\mathcal{A}$ is empty, the set of rigid inclusions is, as in the previous case,

$$\mathcal{R} = \{Fish \sqsubseteq \forall livesIn.Water\}$$

and the set of typicality properties $\mathcal{T}$ is now as follows:

(1) 0.9 :: $\mathbf{T}(Fish) \sqsubseteq \neg Affectionate$
(2) 0.8 :: $\mathbf{T}(Fish) \sqsubseteq Greyish$
(3) 0.8 :: $\mathbf{T}(Fish) \sqsubseteq Scaly$
(4) 0.9 :: $\mathbf{T}(Fish) \sqsubseteq \neg Warm$
(5) 0.8 :: $\mathbf{T}(Pet) \sqsubseteq \forall livesIn.(\neg Water)$
(6) 0.9 :: $\mathbf{T}(Pet) \sqsubseteq Affectionate$
(7) 0.95 :: $\mathbf{T}(Pet) \sqsubseteq Warm$

As in the previous case, by the properties of the typicality operator $\mathbf{T}$, we inherit from the underlying $\mathcal{ALC} + \mathbf{T_R}$ that $\mathbf{T}(Pet \sqcap Fish) \sqsubseteq \forall livesIn.Water$.

In the novel set of generated scenarios we have that the scenario with the highest probability (up to 35%) is both trivial and inconsistent, therefore the following blocks of scenarios are considered.

Scenarios with probabilities from $35,459\%$ down to 0.117% are inconsistent. Considering blocks of consistent scenarios in descending order of their probabilities, all scenarios with probabilities ranging from 0.109% down to 0.027% are discarded since they are either trivial, i.e. all the properties of the HEAD are inherited, or they privilege the MODIFIER against the HEAD.

The first valid block, i.e. block # 4 in Figure 2, with probability 0.013%, contains four consistent scenarios. Again, two of them are discarded for HEAD/MODIFIER heuristics. Among the remaining two valid scenarios, only one of them is able to model the pet fish phenomenon: it includes three out of four properties of the HEAD, therefore it is not trivial and it is selected by the logic $\mathbf{T}^{\mathsf{CL}}$ for the composition of the two initial prototypes.

In this case $\mathbf{T}^{\mathsf{CL}}$ remains agnostic with respect to the final selection of the scenario. For the sake of our purpose, indeed, the fact that $\mathbf{T}^{\mathsf{CL}}$ selects the block containing the scenario with the right answer to the PET FISH problem is, even if a non optimal solution, a satisfying one. In fact, the assigned probabilities in the initial KB can be considered a more complicated testbed for this combination with respect to the previous setup. Indeed, in the latter case, there is i) no probability in the HEAD higher than those in the MODIFIER and ii) the attribute that is expected to be discarded has not a lower probability than the other properties that are assumed to be inherited by the compound concept.

In this respect, it seems an acceptable compromise that the correct solution is included in the set of selected scenarios having the same probability[13]. It is worth-noticing that also in the PET FISH formalization proposed by [Lewis and Lawry (2016)] it is explicitly mentioned that the result of the right composition in the PET FISH problem is reliant on the values chosen for the prototypes of PET and FISH (as this example shows). In particular, with respect to the proposal by [Lewis and Lawry (2016)], where the authors assume that the concepts of PET and FISH are exclusively represented by the properties: "lives in house" and "furry" (with weighed values of 0.6 and 0.4 respectively) and "lives in water" and "scaly" respectively (both with values 0.5), our logic is able to handle a more complicated running example and is able to deal with the problem of blocking - under reasonable circumstances (including some

[13] As we will see in the next sections, in application domains in which $\mathbf{T}^{\mathsf{CL}}$ is exploited, e.g. in the field of computational creativity, where it is not always easy to define which is the "correct" combination, this situation is not uncommon. As a consequence, in such cases the final decision about what selection represents the most appropriate combination, is left to the human decision makers.

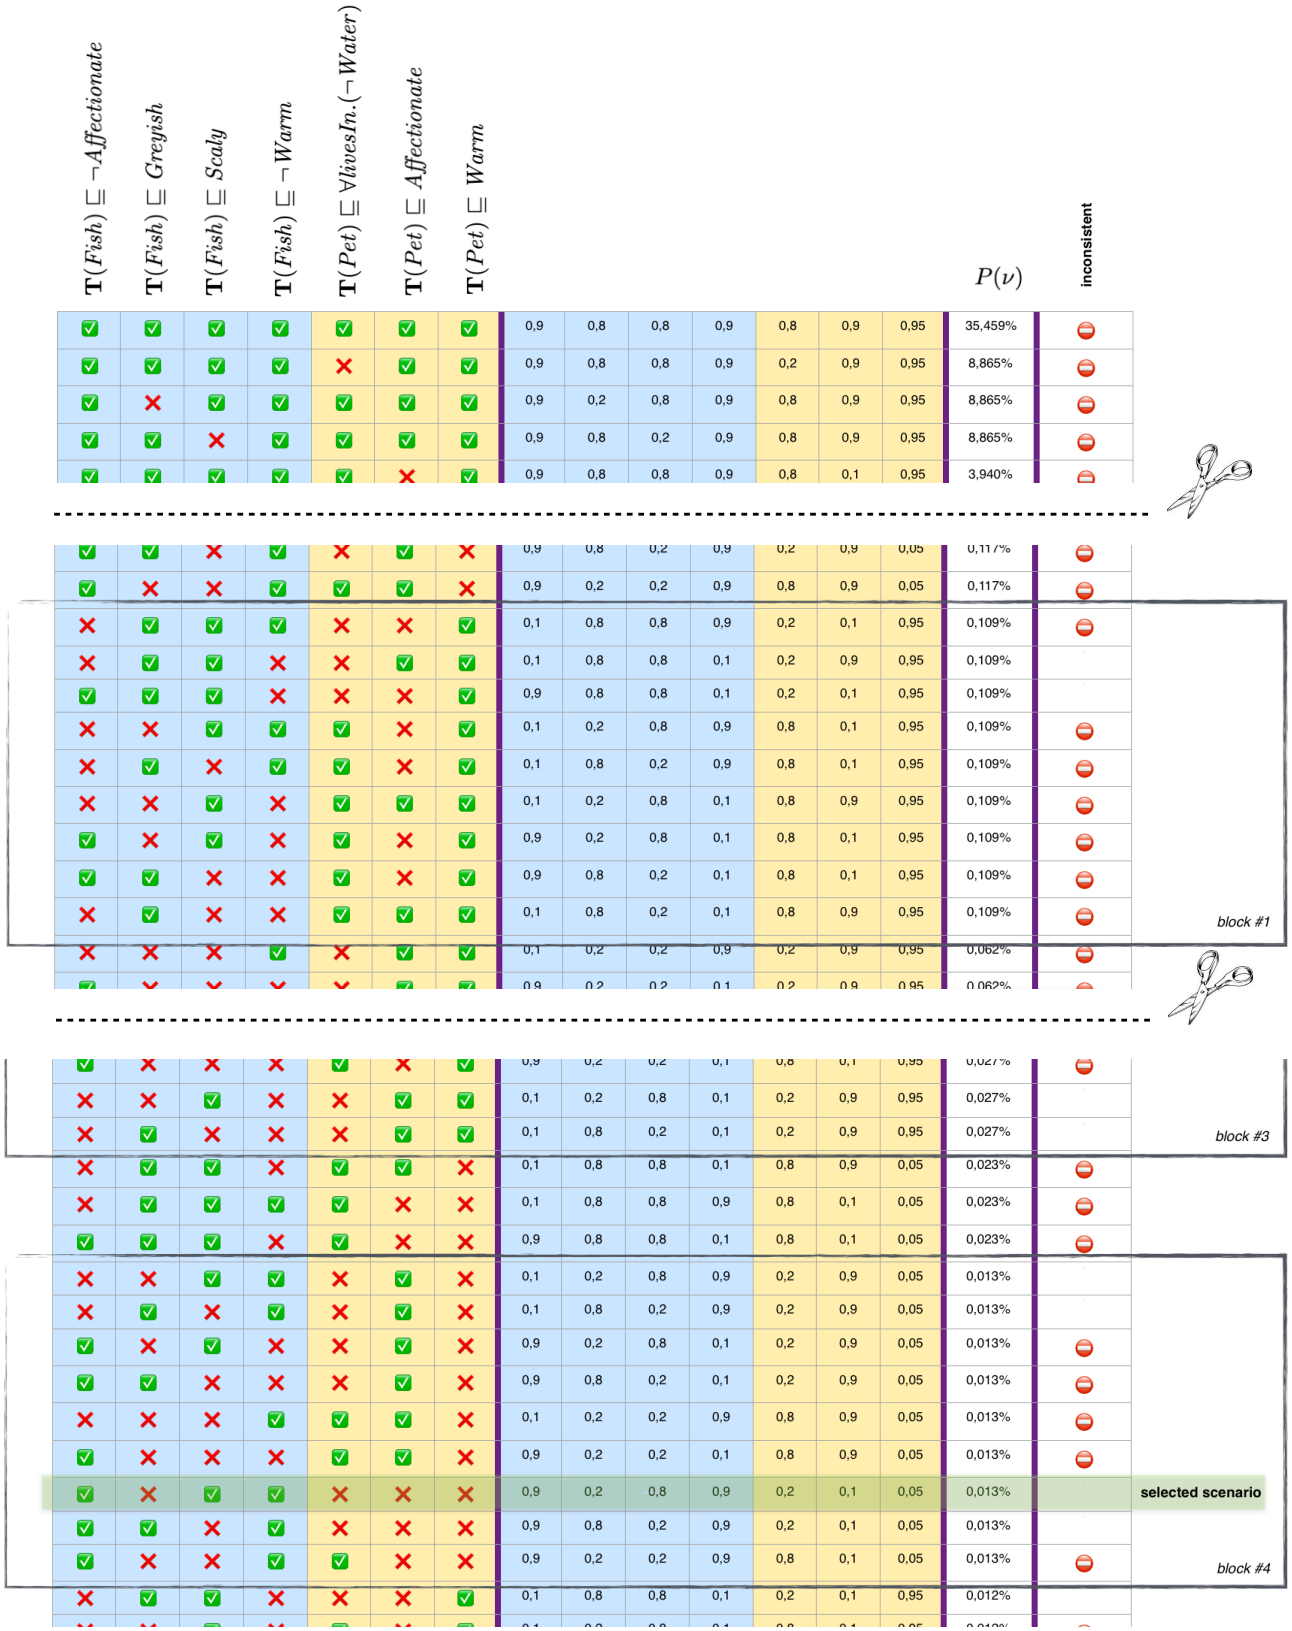

Figure 2.: Scenarios of the combination of Pet and $Fish$. For the sake of readability we cut off the first part of the figure since the first blocks of the combination, i.e. 43 scenarios, are all inconsistent.

of the most challenging ones) - the inheritance of undesired properties (e.g like in the case of "greyish", not presented in the formalization of [Lewis and Lawry (2016)] but available in all the descriptions of the PET FISH problem provided by the cognitive science literature on the phenomenon). In the case of $\mathbf{T}^{\mathsf{CL}}$, the only situation that cannot be handled is represented by a starting KB assigning the highest degree of belief/probability value to the property of the combination that should be discarded. To the best of our knowledge, however, there is no formalization that is currently able to handle such a situation.

4.2. Linda the feminist bank teller

We now exploit the logic $\mathbf{T}^{\mathsf{CL}}$ in order to tackle the conjunction fallacy problem (or "Linda Problem"). The problem configuration is as follows: let us suppose to know that Linda is a 31 years old, single, outspoken, and bright lady. She majored in philosophy and was concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. When asked to rank the probability of the statements 1) "Linda is a bank teller" and 2) "Linda is a bank teller and is active in the feminist movement", the majority of people rank 2) as more probable than 1), violating the classic probability rules. In our logic, let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB, where $\mathcal{A} = \emptyset$, $\mathcal{T}$ is:

$0.8 :: \mathbf{T}(\mathit{Feminist}) \sqsubseteq \mathit{Bright}$
$0.9 :: \mathbf{T}(\mathit{Feminist}) \sqsubseteq \mathit{OutSpoken}$
$0.8 :: \mathbf{T}(\mathit{Feminist}) \sqsubseteq \exists \mathit{fightsFor}.\mathit{SocialJustice}$
$0.9 :: \mathbf{T}(\mathit{Feminist}) \sqsubseteq \mathit{Environmentalist}$
$0.6 :: \mathbf{T}(\mathit{BankTeller}) \sqsubseteq \neg\exists \mathit{fightsFor}.\mathit{SocialJustice}$
$0.8 :: \mathbf{T}(\mathit{BankTeller}) \sqsubseteq \mathit{Calm}$

and $\mathcal{R}$ is as follows:

$\mathit{BankTeller} \sqsubseteq \exists \mathit{isEmployed}.\mathit{Bank}$
$\mathit{Feminist} \sqsubseteq \exists \mathit{believesIn}.\mathit{Feminism}$
$\mathit{Feminist} \sqsubseteq \mathit{Female}$
$\mathit{Environmentalist} \sqsubseteq \exists \mathit{isAgainst}.\mathit{NuclearEnergyDevelopment}$

Let us consider the compound concept $\mathit{Feminist} \sqcap \mathit{BankTeller}$. It can be obtained in two different ways, namely by choosing $\mathit{Feminist}$ as the HEAD and $\mathit{BankTeller}$ as the MODIFIER, or vice versa. First, we consider $\mathit{Feminist}$ as the HEAD. In $\mathbf{T}^{\mathsf{CL}}$, the compound concept inherits all the rigid properties, that is to say $\mathit{Feminist} \sqcap \mathit{BankTeller}$ is included in $\exists \mathit{isEmployed}.\mathit{Bank}$, in $\exists \mathit{believesIn}.\mathit{Feminism}$ and in Female. Concerning the typical properties, two of them are in contrast, namely typical feminists fight for social justice, whereas typical bank tellers do not. All the scenarios including both $\mathbf{T}(\mathit{Feminist}) \sqsubseteq \exists \mathit{fightsFor}.\mathit{SocialJustice}$ and $\mathbf{T}(\mathit{BankTeller}) \sqsubseteq \neg\exists \mathit{fightsFor}.\mathit{SocialJustice}$ are then inconsistent. Concerning the remaining ones, scenarios including $\mathbf{T}(\mathit{BankTeller}) \sqsubseteq \neg\exists \mathit{fightsFor}.\mathit{SocialJustice}$ are discarded, in favor of scenarios including $\mathbf{T}(\mathit{Feminist}) \sqsubseteq \exists \mathit{fightsFor}.\mathit{SocialJustice}$. The most obvious scenario, with the highest probability, corresponds to the one including all the typicality inclusions related to the HEAD. In the logic $\mathbf{T}^{\mathsf{CL}}$ we discard it and we focus on the remaining ones. Among them, one of the scenarios having the highest probability is the one not including $\mathbf{T}(\mathit{Feminist}) \sqsubseteq \mathit{Bright}$. This scenario defines the following $\mathit{Feminist} \sqcap \mathit{BankTeller}$-revised knowledge base:

$0.9 :: \mathbf{T}(\mathit{Feminist} \sqcap \mathit{BankTeller}) \sqsubseteq \mathit{OutSpoken}$
$0.8 :: \mathbf{T}(\mathit{Feminist} \sqcap \mathit{BankTeller}) \sqsubseteq \exists \mathit{fightsFor}.\mathit{SocialJustice}$
$0.9 :: \mathbf{T}(\mathit{Feminist} \sqcap \mathit{BankTeller}) \sqsubseteq \mathit{Environmentalist}$
$0.8 :: \mathbf{T}(\mathit{Feminist} \sqcap \mathit{BankTeller}) \sqsubseteq \mathit{Calm}$

Let us now consider the case of the instance Linda, that is described as follows:

$YoungWoman(linda)$
$\exists graduatedIn.Philosophy(linda)$
$Outspoken(linda)$
$Bright(linda)$
$Single(linda)$
$\exists fightsFor.SocialJustice(linda)$
$\exists isAgainst.NuclearEnergyDevelopment(linda)$

In our logic, solving the conjunction fallacy problem [14] means that we have to find the most appropriate category for Linda. In our case the choice is between $BankTeller$ and $Feminist \sqcap BankTeller$. We can assume that, in absence of any other information, the described properties that are explicitly assigned to the instance Linda can be set to a default probability value of 0.6 (that is to say: the asserted properties about Linda are considerer "typical enough" for her description). Let us first consider the $Feminist \sqcap BankTeller$-revised knowledge base, with an ABox asserting that Linda is a bank teller, that is to say

$$\mathcal{A}_1 = \{BankTeller(linda)\},$$

and let us consider each property of the instance Linda and the associated probability of entailment. Observe that none of such properties are entailed by the $Feminist \sqcap BankTeller$-revised knowledge base with $\mathcal{A}_1$, therefore, for each property of the form $D(linda)$ we have that $\mathbb{P}(D(linda), 0.6) = 0$ (by Definition 3.9). On the other hand, let us consider an ABox asserting that Linda is a feminist bank teller, namely

$$\mathcal{A}_2 = \{(Feminist \sqcap BankTeller)(linda)\}.$$

In this case, we have that:

- $YoungWoman(linda)$ is not entailed from the $Feminist \sqcap BankTeller$-revised knowledge base, therefore
  $\mathbb{P}(YoungWoman(linda), 0.6) = 0$; the same for $\exists graduatedIn.Philosophy(linda)$ and $Single(linda)$;
- $Outspoken(linda)$ is entailed from the $Feminist \sqcap BankTeller$-revised knowledge base with $\mathcal{A}_2$, then, by Definition 3.9, we have $\mathbb{P}(Outspoken(linda), 0.6) = 0.6 \times 0.9 = 0.54$, where 0.9 is the probability of $\mathbf{T}(Feminist \sqcap BankTeller) \sqsubseteq OutSpoken$ in the $Feminist \sqcap BankTeller$-revised KB;
- the same holds for $\exists fightsFor.SocialJustice\ (linda)$, which is entailed by using $\mathcal{A}_2$: in this case, we have that
  $\mathbb{P}(\exists fightsFor.SocialJustice(linda), 0.6) = 0.6 \times 0.8 = 0.48$;
- the fact $\exists isAgainst.NuclearEnergyDevelopment(linda)$ is entailed by using $\mathcal{A}_2$. Observe that $Environmentalist \sqsubseteq NuclearEnergyDevelopment$ follows from $\mathcal{R}$ in standard $\mathcal{ALC}$, then $\mathbb{P}(\exists isAgainst.NuclearEnergyDevelopment\ (linda), 0.6) =$

[14] The attempt of modelling a reasoning error could seem *prima facie*, counterintuitive in an AI setting. However, it is worth-noting that this type of reasoning corresponds to a very powerful evolutionary heuristics developed by humans and strongly relying on common-sense knowledge. The use of typical knowledge in cognitive tasks, in fact, has to do with the constraints that concern every finite agent that has a limited access to the knowledge relevant for a given task. Consider for example the following variant of the Linda problem. Let us suppose that a certain individual Pluto is described as follows. He weighs about 250 kg, and he is approximately two meters tall. His body is covered with a thick, dark fur, he has a large mouth with robust teeth and paws with long claws. He roars and growls. Now, given this information, we have to evaluate the plausibility of the two following alternatives: a) Pippo is a mammal; b) Pippo is a mammal, and he is wild and dangerous. Which is the "correct" answer? According to the dictates of the normative theory of probability, it is surely a). But if you encounter Pippo in the wilderness, it would probably be best to run.

$0.6 \times 0.9 = 0.54$ by Definition 3.9, where 0.9 is the probability of $\mathbf{T}(\mathit{Feminist} \sqcap \mathit{BankTeller}) \sqsubseteq \mathit{Environmentalist}$ in the $\mathit{Feminist} \sqcap \mathit{BankTeller}$-revised KB.

Computing the sum of the probabilities of the queries of all facts about Linda, we obtain $0.54+0.48+0.54 = 1.56$, to witness that the choice of $\mathcal{A}_2$ is more appropriate w.r.t. the choice of $\mathcal{A}_1$ where the sum is 0. This means that, in our logic, the human choice of classifying Linda as a feminist bank teller sounds perfectly plausible and has to be preferred to the alternative one of classifying her as a bank teller.

Let us now consider the case in which $\mathit{BankTeller}$ is the HEAD. In this case, the $\mathit{BankTeller} \sqcap \mathit{Feminist}$-revised knowledge base would be as follows:

$0.8 :: \mathbf{T}(\mathit{BankTeller} \sqcap \mathit{Feminist}) \sqsubseteq \mathit{Bright}$
$0.9 :: \mathbf{T}(\mathit{BankTeller} \sqcap \mathit{Feminist}) \sqsubseteq \mathit{OutSpoken}$
$0.8 :: \mathbf{T}(\mathit{BankTeller} \sqcap \mathit{Feminist}) \sqsubseteq \exists \mathit{fightsFor}.\mathit{SocialJustice}$
$0.9 :: \mathbf{T}(\mathit{BankTeller} \sqcap \mathit{Feminist}) \sqsubseteq \mathit{Environmentalist}$
$0.8 :: \mathbf{T}(\mathit{BankTeller} \sqcap \mathit{Feminist}) \sqsubseteq \mathit{Calm}$

Also in this case, the probability that Linda is a bank teller feminist is higher (2.04) than the probability of classifying her as a bank teller (as in the previous case, equals to zero).

4.3. Metaphorical Concept Combination: a prototype of the Stone Lion

In this section we consider a particular type of concept combination based on the nonmononotonic character of $\mathbf{T}^{\mathsf{CL}}$. In particular, we take into account a classical case of metaphorical concept composition considered in the field of cognitive semantics: the Stone Lion example [Franks (1995); Gärdenfors (1998, 2014)]. If we consider the concept *Lion* in isolation, typically it is inferred that it is alive, it has fur and a tail, and so on. If we consider the combination of concept *Stone* and *Lion*, on the other hand, the only inherited aspect which is Lion-like is its shape, this means that, in this case, the effect of the combination is obtained because a stone object is metaphorically seen as a lion, due to its shape. Let us consider in detail this example, and let us exploit the logic $\mathbf{T}^{\mathsf{CL}}$ in order to provide a commonsense description of a prototype of the Stone Lion. Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB, where: $\mathcal{A}$ is empty, $\mathcal{R} = \{\mathit{MainColorYellowish} \sqcap \mathit{MainColorGreyish} \sqsubseteq \bot\}$ are empty, and $\mathcal{T}$ is as follows:

(1) $0.9 :: \mathbf{T}(\mathit{Stone}) \sqsubseteq \mathit{HardMaterial}$
(2) $0.8 :: \mathbf{T}(\mathit{Stone}) \sqsubseteq \mathit{MainColorGreyish}$
(3) $0.7 :: \mathbf{T}(\mathit{Stone}) \sqsubseteq \mathit{Rolling}$
(4) $0.8 :: \mathbf{T}(\mathit{Lion}) \sqsubseteq \mathit{MainColorYellowish}$
(5) $0.7 :: \mathbf{T}(\mathit{Lion}) \sqsubseteq \exists \mathit{has}.\mathit{Tail}$

We consider Stone as the HEAD and Lion as the MODIFIER in defining this combination. Figure 3 shows the $2^5 = 32$ different scenarios, one row for each scenario, in descending order of provability. As in the example of the pet fish, inconsistent scenarios are highlighted in the last column.

From a probabilistic perspective, the first three scenarios are inconsistent since both inclusions 2 and 4 are included, in other words we would have $\mathbf{T}(\mathit{Stone} \sqcap \mathit{Lion}) \sqsubseteq \mathit{MainColorGreyish}$ and $\mathbf{T}(\mathit{Stone} \sqcap \mathit{Lion}) \sqsubseteq \mathit{MainColorYellowish}$, having that $\mathit{MainColorGreyish} \sqcap \mathit{MainColorYellowish} \sqsubseteq \bot$. The first block to consider contains two scenarios with probability 7.056%: the first one is discarded since it privileges a property of the MODIFIER (having $\mathit{MainColorYellowish}$) rather than a contrasting one of the HEAD (having $\mathit{MainColorGreyish}$). The second one is discarded, since it allows to inherit all the properties of the HEAD and its then considered as trivial. The following useful block is de-

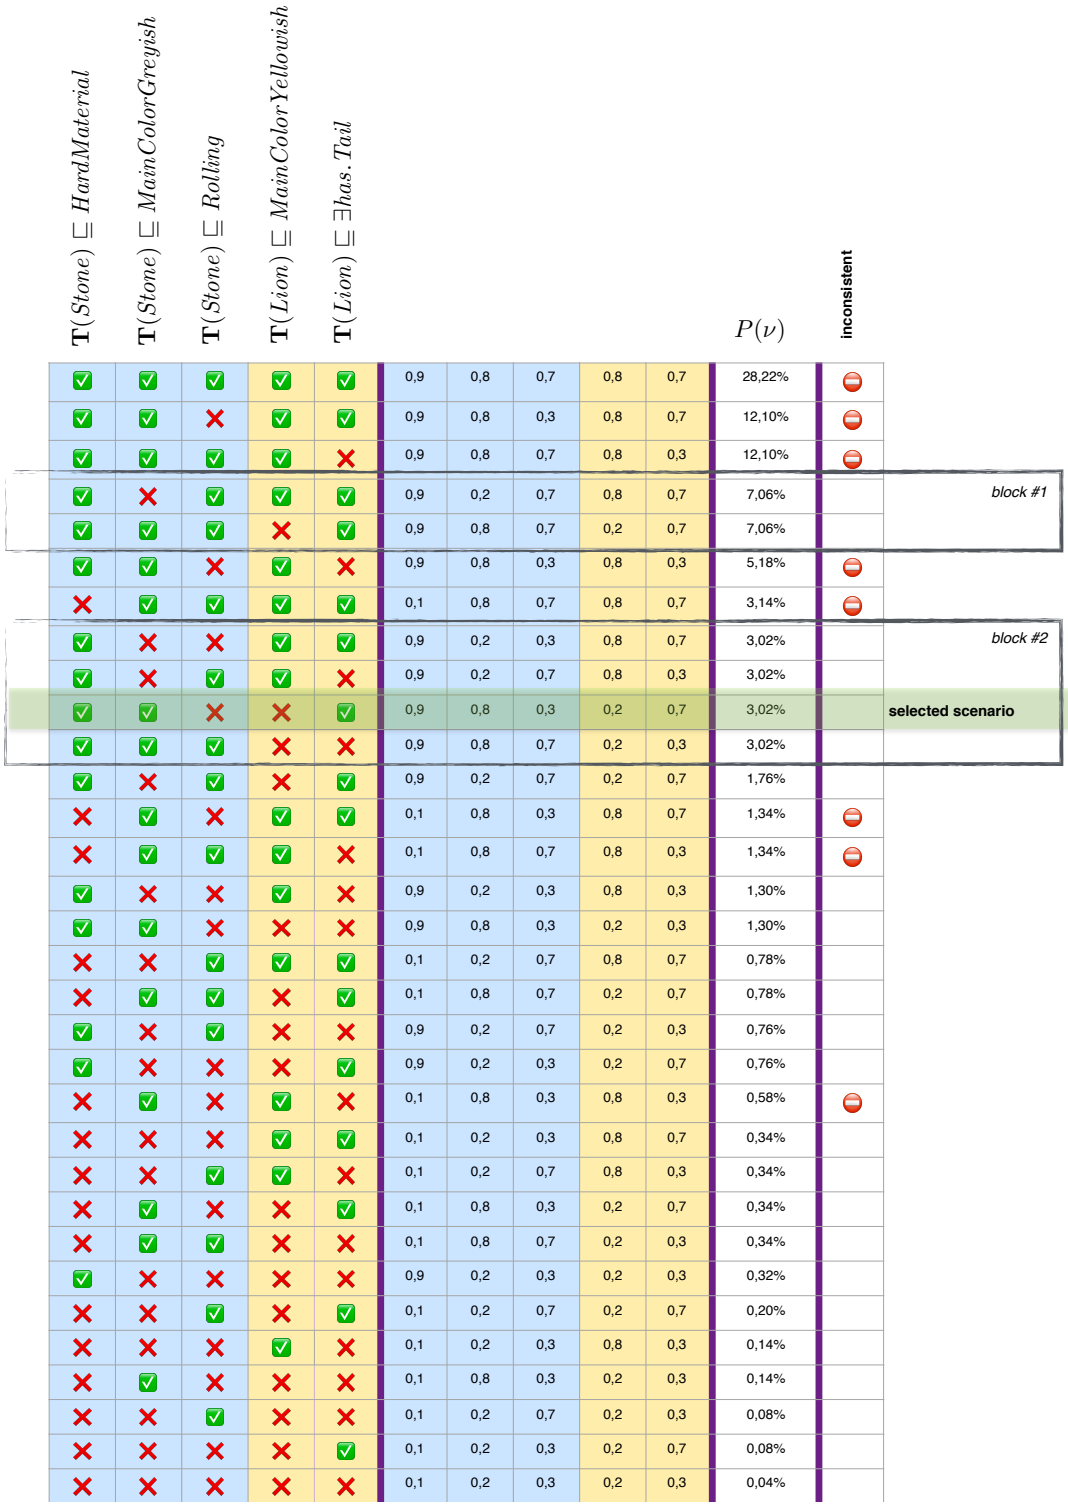

Figure 3.: Scenarios of the combination of $Stone$ and $Lion$.

fined by four scenarios having probability 3.02%:

- the first and the second ones in Figure 3 are discarded since, again, the property $MainColorYellowish$ of the MODIFIER is preferred to the contrasting one of the HEAD;
- the last one is discarded since it contains all the properties of the HEAD, then it is trivial.

The remaining scenario, the third one of the block in Figure 3, is the one selected by the logic $\mathbf{T}^{\mathsf{CL}}$. It contains inclusions 1 and 2 from the HEAD and inclusion 5 from the modifier, obtaining the following $Stone \sqcap Lion$-revised knowledge base:

1. 0.9 :: $\mathbf{T}(Stone \sqcap Lion) \sqsubseteq HardMaterial$
2. 0.8 :: $\mathbf{T}(Stone \sqcap Lion) \sqsubseteq MainColorGreyish$
5. 0.7 :: $\mathbf{T}(Stone \sqcap Lion) \sqsubseteq \exists has.Tail$

Notice that, if $\mathcal{R} = \{Stone \sqsubseteq \neg Breath\}$ and $\mathcal{T}$ also contains $\mathbf{T}(Lion) \sqsubseteq Breath$, in our logic we would also infer that $Stone \sqcap Lion \sqsubseteq \neg Breath$.

5. Artificial Prototypes Composition and Concept Invention

In this section we exploit the logic $\mathbf{T}^{\mathsf{CL}}$ to show both i) how it allows to automatically generate novel, plausible, prototypical concepts by composing two initial prototypes and ii) how it can be used as a generative tool in the field of computational creativity (with applications in the so called creative industry). In detail, we extend preliminary results presented by [Lieto and Pozzato (2019)] and we first show how our logic can model the generation of a quite complex concept recently introduced in the field of narratology, i.e. that one of the ANTI-HERO (a role invented by narratologists to generate new story lines), by combining the typical properties of the concepts HERO and VILLAIN. Of course, the specific domain of the example is not relevant here; our goal is showing how $\mathbf{T}^{\mathsf{CL}}$ can model this kind of prototypical concept composition (a crucial aspect of human concept invention) that, on the other hand, has been proven to be problematic for other kinds of logics (e.g fuzzy logic, [Hampton (2011); Osherson and Smith (1981))]. We then show how the same machinery can be used as a creativity support tool to generate a new type of villain for a video game or a movie.

5.1. Anti Hero

We will take into account the concepts of HERO, ANTI-HERO and VILLAIN extracted by the common sense descriptions coming from the TvTropes repository [15]. In such online repository, typical descriptions of character roles are provided. They can be useful for practitioners of the narrative field in order to design their own character according to the main assets presented in such schemas. In particular, Tropes can be seen as devices and conventions that a writer can reasonably rely on as being present in the audience members' minds and expectations. Regarding the HERO, TvTropes identifies the following relevant representative features: e.g. the fact that it is characterized by his/her fights against the VILLAIN of a story, the fact that his/her actions are necessarily guided by general goals to be achieved in the interest of the collectivity, the fact that they fight against the VILLAIN in a fair way and so on. Examples of such Trope are: Superman, Flash Gordon etc.. The ANTI-HERO, on the other hand, is described as characterized by the fact of sharing most of its typical traits with the

[15] `https://tvtropes.org`

HERO (e.g. the fact that it is the protagonist of a plot fighting against the VILLAIN of the story); however, his/her moves are not guided by a general spirit of sacrifice for the collectivity but, rather, they are usually based on some personal motivations that, incidentally and/or indirectly, coincide with the needs of the collectivity. Furthermore the ANTI-HERO may also act in a not fair way in order to achieve the desired goal. A classical example of such trope is Batman, whose moves are guided by his desire of revenge. Finally the VILLAIN is represented as a classic negative role in a plot and is characterized as the main opponent of the protagonist/HERO. In addition to this classical contraposition, TvTropes also reports some physical elements characterizing such role from a visual point of view. For example: the characters of this Trope are usually physically endowed with some demoniac cues (e.g. they have the "eyes of fire"). Finally, they are guided by negative moral values. Examples of such role can be easily taken from the classical literature to the modern comics. Some representative exemplars are Cruella de Vil in Disney's filmic saga or Voldemort in Harry Potter.

Let us now exploit our logic $\mathbf{T}^{\mathsf{CL}}$ in order to define a prototype of ANTI-HERO. First of all, we define a knowledge base describing both rigid and typical properties of concepts HERO and VILLAIN, then we rely on the logic $\mathbf{T}^{\mathsf{CL}}$ in order to formalize an $AntiHero$-revised knowledge base.

Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB, where the ABox $\mathcal{A}$ is empty. Concerning rigid properties, let $\mathcal{R}$ be as follows:

R1 $Hero \sqsubseteq \exists hasOpponent.Villain$
R2 $Villain \sqsubseteq \forall fightsFor.PersonalGoal$
R3 $Villain \sqsubseteq WithNegativeMoralValues$
R4 $CollectiveGoal \sqcap PersonalGoal \sqsubseteq \bot$
R5 $WithPositiveMoralValues \sqcap WithNegativeMoralValues \sqsubseteq \bot$
R6 $AngelicIconicity \sqcap DemoniacIconicity \sqsubseteq \bot$

Prototypical properties of villains and heroes are described in $\mathcal{T}$ is as follows:

T1 $0.95 :: \mathbf{T}(Hero) \sqsubseteq Protagonist$
T2 $0.85 :: \mathbf{T}(Hero) \sqsubseteq \exists fightsFor.CollectiveGoal$
T3 $0.9 :: \mathbf{T}(Hero) \sqsubseteq WithPositiveMoralValues$
T4 $0.6 :: \mathbf{T}(Hero) \sqsubseteq AngelicIconicity$
T5 $0.75 :: \mathbf{T}(Villain) \sqsubseteq DemoniacIconicity$
T6 $0.8 :: \mathbf{T}(Villain) \sqsubseteq Implulsive$
T7 $0.75 :: \mathbf{T}(Villain) \sqsubseteq Protagonist$

We make use of the logic $\mathbf{T}^{\mathsf{CL}}$ in order to build the compound concept $AntiHero$ as the result of the combination of concepts $Hero$ and $Villain$. Differently from what the natural language seems to suggest, we consider this compound concept by assuming that the HEAD is $Villain$ (since the ANTI-HERO shares more typical traits with this concept than with the HERO concept).

First of all, we have that the compound concepts inherits all the rigid properties of both its components (if not contradictory), therefore in the logic $\mathbf{T}^{\mathsf{CL}}$ we have that:

(i) $AntiHero \sqsubseteq \exists hasOpponent.Villain$
(ii) $AntiHero \sqsubseteq \forall fightsFor.PersonalGoal$
(iii) $AntiHero \sqsubseteq WithNegativeMoralValues$

For the typical properties, we consider all the $2^7 = 256$ different scenarios obtained from all possible selections about inclusion in $\mathcal{T}$. Some of them are inconsistent, namely those including either axiom T2 or axiom T3, since they would ascribe properties in contrast with inherited rigid properties of (ii) and (iii): rigid properties impose that an anti hero has negative

moral values, and all his goals are personal, therefore he is an atypical hero in those respects (T2 states that typical heroes fights also for some collective goals, whereas T3 states that normally heroes have positive moral values). Also scenarios containing both axioms T4 and T5 are inconsistent, since the fact that the concepts $AngelicIconicity$ and $DemoniacIconicity$ are disjoint (formalized by R6).

Let us consider the remaining, consistent scenarios: the one having the highest probability considers all the properties of both concepts by excluding only $AngelicIconicity$, that is to say the one with the lowest probability between the two properties in conflict. In $\mathbf{T}^{\mathsf{CL}}$ this scenario is discarded since it is the most trivial one. When we consider scenarios less trivial, i.e., more surprising scenarios (we analyze scenarios in decreasing order of probability), we discard the scenario with probability 0.13%, which includes T4, associated to the MODIFIER, rather than T5, associated to the HEAD, allowing to conclude, in a counter intuitive way, that typical anti heroes have an angelic iconicity rather than a demoniac one.

Next scenarios, sharing the same probability (0.09%), are as follows:

T1 0.95 :: $\mathbf{T}(Hero) \sqsubseteq Protagonist$	T1 0.95 :: $\mathbf{T}(Hero) \sqsubseteq Protagonist$
T5 0.75 :: $\mathbf{T}(Villain) \sqsubseteq DemoniacIconicity$	T6 0.8 :: $\mathbf{T}(Villain) \sqsubseteq Implulsive$
T6 0.8 :: $\mathbf{T}(Villain) \sqsubseteq Implulsive$	T7 0.75 :: $\mathbf{T}(Villain) \sqsubseteq Protagonist$

According to the logic $\mathbf{T}^{\mathsf{CL}}$, both are adequate and represent the outcome of the whole heuristic procedures adopted in $\mathbf{T}^{\mathsf{CL}}$. Probably, in this case, it could be more useful to opt for the solution on the left allowing to inherit a further property (i.e. $DemoniacIconicity$) for the generated prototypical Anti-Hero. However, we remain agnostic about the selection of the final options provided by $\mathbf{T}^{\mathsf{CL}}$. This choice can be plausibly left to human decision makers and based on their own goals.

5.2. Generating a Novel Character via $\mathbf{T}^{\mathsf{CL}}$: a Villain Chair

Let us now exploit our logic $\mathbf{T}^{\mathsf{CL}}$ in order to create a new compound concept: e.g. a new type of villain for a video game or a movie, obtained by as the combination of concepts $Villain$ (as HEAD) and $Chair$ (as MODIFIER). Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB, where: $\mathcal{A}$ is empty, $\mathcal{R}$ is as follows:

R1 $Villain \sqsubseteq \exists fightsFor.PersonalGoal$
R2 $Villain \sqsubseteq Animate$
R3 $Villain \sqsubseteq WithNegativeMoralValues$
R4 $Chair \sqsubseteq \exists hasComponent.SupportingSeatComponent$
R5 $Chair \sqsubseteq \exists hasComponent.Seat$
R6 $CollectiveGoal \sqcap PersonalGoal \sqsubseteq \bot$

and $\mathcal{T}$ is as follows:

T1 0.9 :: $\mathbf{T}(Villain) \sqsubseteq DemoniacIconicity$
T2 0.75 :: $\mathbf{T}(Villain) \sqsubseteq \exists hasOpponent.Hero$
T3 0.75 :: $\mathbf{T}(Villain) \sqsubseteq Protagonist$
T4 0.8 :: $\mathbf{T}(Villain) \sqsubseteq Impulsive$
T5 0.95 :: $\mathbf{T}(Chair) \sqsubseteq \neg Animate$
T6 0.95 :: $\mathbf{T}(Chair) \sqsubseteq \exists hasComponent.Back$
T7 0.65 :: $\mathbf{T}(Chair) \sqsubseteq \exists madeOf.Wood$

T8 0.8 :: $\mathbf{T}(Chair) \sqsubseteq Comfortable$
T9 0.7 :: $\mathbf{T}(Chair) \sqsubseteq Inflammable$

We consider the 512 scenarios, from which we discard the inconsistent ones, namely those including T5: indeed, since R2 imposes that villains are animate, in the underlying $\mathcal{ALC} + \mathbf{T_R}$ we conclude that $Villain \sqcap Chair \sqsubseteq Animate$, therefore all scenarios including T5, imposing that $Villain \sqcap Chair \sqsubseteq \neg Animate$ are inconsistent. We also discard the most obvious scenario including all the typicality inclusions of $\mathcal{R}$, having probability of 14%. We also discard the following trivial scenarios containing all the inclusions related to the HEAD, namely T1, T2, T3, and T4.

The first suitable block according to Algorithm 1 is the one whose scenarios have probability 4.67% and contain all properties coming from the MODIFIER and three out of four properties coming from the HEAD. Such scenarios, defining two alternative revised knowledge bases (one containing T2 and not T3, the other one containing T3 and not T2), are as follows:

T1 0.9 :: $\mathbf{T}(Villain) \sqsubseteq DemoniacIconicity$	T1 0.9 :: $\mathbf{T}(Villain) \sqsubseteq DemoniacIconicity$
T3 0.75 :: $\mathbf{T}(Villain) \sqsubseteq Protagonist$	T2 0.75 :: $\mathbf{T}(Villain) \sqsubseteq \exists hasOpponent.Hero$
T4 0.8 :: $\mathbf{T}(Villain) \sqsubseteq Impulsive$	T4 0.8 :: $\mathbf{T}(Villain) \sqsubseteq Impulsive$
T6 0.95 :: $\mathbf{T}(Chair) \sqsubseteq \exists hasComponent.Back$	T6 0.95 :: $\mathbf{T}(Chair) \sqsubseteq \exists hasComponent.Back$
T7 0.65 :: $\mathbf{T}(Chair) \sqsubseteq \exists madeOf.Wood$	T7 0.65 :: $\mathbf{T}(Chair) \sqsubseteq \exists madeOf.Wood$
T8 0.8 :: $\mathbf{T}(Chair) \sqsubseteq Comfortable$	T8 0.8 :: $\mathbf{T}(Chair) \sqsubseteq Comfortable$
T9 0.7 :: $\mathbf{T}(Chair) \sqsubseteq Inflammable$	T9 0.7 :: $\mathbf{T}(Chair) \sqsubseteq Inflammable$

According to Algorithm 1, both proposals neither are trivial (not all properties of the HEAD are ascribed to the combined concept) nor they give preference to the MODIFIER with respect to the HEAD for conflicting typical properties (in the example, we have only a conflict between a rigid and typical property, the above mentioned R2 and T5). These scenarios are the preferred ones selected by the logic $\mathbf{T}^{\mathsf{CL}}$.

However, in this application setting, we could imagine to use our framework as a creativity support tool and thus considering alternative - more surprising - scenarios by adding additional constraints. For example, we could impose that the compound concept should inherit exactly six properties. In this case, we would get that the scenario having the highest probability (3.2%) is the one including all the properties of the HEAD, namely T1, T2, T3 and T4, and two out of four properties of the MODIFIER, namely T6 and T8. Due to its triviality, this scenario is discarded, in favor of the following block of two scenarios (probability 2.51%), obtained by excluding T7 of the MODIFIER and one out of four properties of the HEAD:

T1 0.9 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq DemoniacIconicity$	T1 0.9 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq DemoniacIconicity$
T3 0.75 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Protagonist$	T2 0.75 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq \exists hasOpponent.Hero$
T4 0.8 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Impulsive$	T4 0.8 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Impulsive$
T6 0.95 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq \exists hasComponent.Back$	T6 0.95 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq \exists hasComponent.Back$
T8 0.8 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Comfortable$	T8 0.8 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Comfortable$
T9 0.7 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Inflammable$	T9 0.7 :: $\mathbf{T}(Villain \sqcap Chair) \sqsubseteq Inflammable$

These are the scenarios selected by $\mathbf{T}^{\mathsf{CL}}$ and are those providing plausible but not obvious, creative definitions of a villain chair.

5.3. Generating Compounds by Combining Multiple Concepts: a prototype for the Chimera

In this section we show how the mechanism of the logic $\mathbf{T}^{\mathsf{CL}}$ can be used in order to combine more than two atomic concepts (a case never taken into account in previous formalizations, e.g. in [Eppe et al. (2018); Lewis and Lawry (2016)] etc.). Such examples show how the mechanisms behind $\mathbf{T}^{\mathsf{CL}}$ can be iteratively applied without loss of efficacy in the produced output. In our opinion, this is a symptom of the fact that our framework is able to actually capture some foundational elements of common-sense conceptual compositionality.

Let us consider the example of a chimera, a mythological hybrid entity composed of the parts of more than one animal. It is usually depicted as a lion, with the head of a goat positioned in the center of its body, and a tail ended with the head of a fire-breathing dragon [16].

First of all, let us describe the three atomic concepts to be combined, namely $Lion$, $Goat$, and $Dragon$. Let $\mathcal{K}$ be as follows:

(1) $Lion \sqsubseteq Animal$
(2) $0.7 :: \mathbf{T}(Lion) \sqsubseteq MainColorYellowish$
(3) $0.9 :: \mathbf{T}(Lion) \sqsubseteq \exists has.Tail$
(4) $0.8 :: \mathbf{T}(Lion) \sqsubseteq \exists has.Mane$
(5) $Goat \sqsubseteq Animal$
(6) $0.7 :: \mathbf{T}(Goat) \sqsubseteq MainColorWhitish$
(7) $0.75 :: \mathbf{T}(Goat) \sqsubseteq \exists has.Tail$
(8) $0.9 :: \mathbf{T}(Goat) \sqsubseteq \exists has.Horns$
(9) $0.9 :: \mathbf{T}(Goat) \sqsubseteq \exists has.Beards$
(10) $Dragon \sqsubseteq LegendaryCreature$
(11) $0.9 :: \mathbf{T}(Dragon) \sqsubseteq FireBreathing$
(12) $0.8 :: \mathbf{T}(Dragon) \sqsubseteq Fly$
(13) $0.9 :: \mathbf{T}(Dragon) \sqsubseteq Aggressive$
(14) $Whitish \sqcap Yellowish \sqsubseteq \bot$

We consider that $Lion$ as the HEAD, whereas both $Goat$ and $Dragon$ are modifiers.

We discard all the inconsistent scenarios including both 2 and 6. We also discard the most trivial scenario including all the typicality inclusions of the HEAD, whose probability is 5.95% and, in the same block of scenarios, we also discard the one giving preference to the MODIFIER $Goat$ with respect to the HEAD. More precisely the scenario preferring the inclusion 6 to 2. Among the remaining ones, the selected scenario having the highest probability (2.55%) is as follows, not including both 2 and 6 ($Chimera$-revised knowledge base):

- $Chimera \sqsubseteq Animal$
- $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Tail$
- $0.8 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Mane$
- $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Horns$
- $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Beards$
- $Chimera \sqsubseteq LegendaryCreature$
- $0.9 :: \mathbf{T}(Chimera) \sqsubseteq FireBreathing$
- $0.8 :: \mathbf{T}(Chimera) \sqsubseteq Fly$

[16] Technically this kind of combination is called conceptual blending and is slightly different from a classical combination since the generated concept is an entirely new one and is not a subset of classes generating it, see [Nagai and Taura (2006)] for more details on the subtle differences between the two tasks.

- $0.9 :: \mathbf{T}(Chimera) \sqsubseteq Aggressive$
- $White \sqcap Yellow \sqsubseteq \bot$

6. Iterated Generation of Concepts: Combining Concepts from C-revised Knowledge Bases

A *C-revised* knowledge base in the logic $\mathbf{T}^{\mathsf{CL}}$ is still in the language of the $\mathbf{T}^{\mathsf{CL}}$ logic. This allows us to iteratively repeat the same procedure in order to combine not only atomic concepts, but also compound concepts. In this section we show that our approach can handle also the concept combination of C-revised knowledge bases.

Let us consider the compound concept obtained as the combination of an anti-hero – as described by the $AntiHero$-knowledge base of Section 5.1[17] – and a chimera – $Chimera$-revised knowledge base in Section 5.3.

The starting TBox $\mathcal{T}$ is as follows:

(1) $Chimera \sqsubseteq Animal$
(2) $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Tail$
(3) $0.8 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Mane$
(4) $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Horns$
(5) $0.9 :: \mathbf{T}(Chimera) \sqsubseteq \exists has.Beards$
(6) $Chimera \sqsubseteq LegendaryCreature$
(7) $0.9 :: \mathbf{T}(Chimera) \sqsubseteq FireBreathing$
(8) $0.8 :: \mathbf{T}(Chimera) \sqsubseteq Fly$
(9) $0.9 :: \mathbf{T}(Chimera) \sqsubseteq Aggressive$
(10) $AntiHero \sqsubseteq \exists hasOpponent.Villain$
(11) $AntiHero \sqsubseteq \exists fightsFor.PersonalGoal$
(12) $AntiHero \sqsubseteq WithNegativeMoralValues$
(13) $0.95 :: \mathbf{T}(AntiHero) \sqsubseteq Protagonist$
(14) $0.75 :: \mathbf{T}(AntiHero) \sqsubseteq DemoniacIconicity$
(15) $0.8 :: \mathbf{T}(AntiHero) \sqsubseteq Impulsive$

We discard the most obvious scenario having a probability of 20.4% and considering all the typicality inclusions. Also the immediately lower scenario with probability 7.18% is discarded since it inherits all the properties of the HEAD (and, therefore, is trivial). As for the case of the villain chair of Section 5, we firstly consider the following HEAD-MODIFIER combination: chimera is the HEAD, whereas anti-hero is the MODIFIER. Given this, the scenarios having the highest probability (5.3%), giving preference to the HEAD, are the ones discarding either 3 (has mane) or 8 (fly), and including all the other typicality assumptions. These scenarios are those preferred from a cognitive point of view.

- $Chimera \sqsubseteq Animal$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Tail$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Horns$
- $Chimera \sqsubseteq LegendaryCreature$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Aggressive$
- $AntiHero \sqsubseteq \exists hasOpponent.Villain$
- $AntiHero \sqsubseteq \exists fightsFor.PersonalGoal$

[17] In this example we consider the first $AntiHero$-revised knowledge base obtained in Section 5.1.

- $AntiHero \sqsubseteq WithNegativeMoralValues$
- $0.75 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Protagonist$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq DemoniacIconicity$
- $0.8 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Impulsive$

- $Chimera \sqsubseteq Animal$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Beards$
- $Chimera \sqsubseteq LegendaryCreature$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq FireBreathing$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Aggressive$
- $AntiHero \sqsubseteq \exists hasOpponent.Villain$
- $AntiHero \sqsubseteq \exists fightsFor.PersonalGoal$
- $AntiHero \sqcap Chimera \sqsubseteq WithNegativeMoralValues$
- $0.75 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Protagonist$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq DemoniacIconicity$
- $0.8 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Impulsive$

We could also consider more surprising scenarios by tuning both the number of properties inherited by the compound concepts and by reverting the assignment of the HEAD-MODIFIER pair among the concepts. In the first case, let us restrict our concern to scenarios allowing the chimera anti-hero to inherit 6 properties[18] from the original concepts. In this case, the scenario selected according to $\mathbf{T}^{\mathsf{CL}}$ has probability 0.11% and includes 5 out of 7 inclusions of the HEAD, namely 2, 4, 5, 7, 9, and only one of the MODIFIER, namely 13. In this case, the $AntiHero - Chimera$-revised knowledge base is as follows:

- $Chimera \sqsubseteq Animal$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Tail$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Horns$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Beards$
- $Chimera \sqsubseteq LegendaryCreature$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq FireBreathing$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Aggressive$
- $AntiHero \sqsubseteq \exists hasOpponent.Villain$
- $AntiHero \sqsubseteq \exists fightsFor.PersonalGoal$
- $AntiHero \sqsubseteq WithNegativeMoralValues$
- $0.95 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Protagonist$

Notice that, as in all the previous examples, the combined concept implicitly inherits all the rigid properties of both HEAD and MODIFIER: here, for instance, we derive that $AntiHero \sqcap Chimera \sqsubseteq Animal$, as well as $AntiHero \sqcap Chimera \sqsubseteq \exists hasOpponent.Villain$, by the rigid properties $Chimera \sqsubseteq Animal$ and $AntiHero \sqsubseteq \exists hasOpponent.Villain$, respectively.

Let us conclude this example by swapping the HEAD and the MODIFIER, namely we exploit the logic $\mathbf{T}^{\mathsf{CL}}$ in order to define a anti-hero chimera by considering the concept $AntiHero$ as the HEAD and the concept $Chimera$ as the MODIFIER. This is just based on the assumption that the role played by the novel character can be more relevant from a narrative point of view. Again, we discard the most obvious scenario having a probability of 20.4% and considering all the typicality inclusions. However, the scenario with immediately lower probability

[18] Of course, the number of properties can be considered as a parameter through which it is possible to play with the mechanisms underlying the logic $\mathbf{T}^{\mathsf{CL}}$.

(7.18%) is selected, since it inherits only 2 out of 3 properties from the HEAD, namely those corresponding to inclusions 13 and 14. The resulting $AntiHero - Chimera$-revised knowledge base is as follows:

- $Chimera \sqsubseteq Animal$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Tail$
- $0.8 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Mane$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Horns$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq \exists has.Beards$
- $Chimera \sqsubseteq LegendaryCreature$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq FireBreathing$
- $0.8 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Fly$
- $0.9 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Aggressive$
- $AntiHero \sqsubseteq \exists hasOpponent.Villain$
- $AntiHero \sqsubseteq \exists fightsFor.PersonalGoal$
- $AntiHero \sqsubseteq WithNegativeMoralValues$
- $0.95 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Protagonist$
- $0.8 :: \mathbf{T}(AntiHero \sqcap Chimera) \sqsubseteq Impulsive$

7. Conclusions, Related Works and Future Research

We have introduced a nonmonotonic Description Logic $\mathbf{T}^{\mathsf{CL}}$ for concept combination, extending the DL of typicality $\mathcal{ALC} + \mathbf{T_R}$ with a DISPONTE semantics and with a blocking inheritance selection heuristics coming from the cognitive semantics. This logic enjoys good computational properties, since entailment in it remains `ExpTime` as the underlying monotonic $\mathcal{ALC}$, and is able to take into account the concept combination of prototypical properties. To this aim, the logic $\mathbf{T}^{\mathsf{CL}}$ allows to have inclusions of the form $p :: \mathbf{T}(C) \sqsubseteq D$, representing that, with a probability p, typical Cs are also Ds. Then, several different scenarios – having different probabilities – are described by including or not such inclusions, and prototypical properties of combinations of concepts are obtained by restricting reasoning services to scenarios having suitable probabilities, excluding "trivial" ones with the highest probabilities.

Several approaches in extending DLs with nonmonotonic capabilities have been proposed in the literature. All these approaches are essentially based on the integration of DLs with well established nonmonotonic reasoning mechanisms [Baader and Hollunder (1995b); Bonatti et al. (2015, 2009); Casini and Straccia (2010, 2013); Donini et al. (2002b); Giordano et al. (2013a)], ranging from Reiter's defaults to minimal knowledge and negation as failure. In [Lukasiewicz (2008)] two probabilistic extensions of Description Logics $\mathcal{SHIF}(\mathbf{D})$ and $\mathcal{SHOIN}(\mathbf{D})$ are introduced. These extensions are semantically based on the notion of probabilistic lexicographic entailment [Lehmann (1995)] and allow to represent and reason about prototypical properties of classes that are semantically interpreted as lexicographic entailment introduced by Lehmann from conditional knowledge bases. Intuitively, the basic idea is to interpret inclusions of the TBox and facts in the ABox as probabilistic knowledge about random and concrete instances of concepts. As an example, in these extensions one can express that "typically, a randomly chosen student makes use of social networks with a probability of at least 70%" with a formula of the form $(SocialNetworkUser \mid Student)[0.7, 1]$. As the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ underlying our work, the lexicographic entailment defined in [Lukasiewicz (2008)] inherits interesting and useful nonmonotonic properties from lexicographic entailment in [Lehmann (1995)], such as specificity, rational monotonicity and some

forms of irrelevance. As a difference with our approach, here probabilities are used in order to capture an alternative notion of typicality of concepts, rather than to express degrees of belief of typicality inclusions that are at the base of our notion of our machinery for concept combination. All these approaches could be seen as possible alternative with respect to the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$, representing one of the building blocks of the logic $\mathbf{T}^{\mathsf{CL}}$. However, we do not investigate here to what extent these approaches can actually be compliant with respect to the problem of handling typicality based concept combinations. To the best of our knowledge, there are not existing works exploiting these logics to handle the problem in focus and we believe this aspect is worth-considering for future works.

Another related work with respect to the proposed formalism is in [Pozzato (2017b)]. Here the author extends the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ by means of probabilities equipping typicality inclusions of the form $\mathbf{T}(C) \sqsubseteq_p D$, whose intuitive meaning is that, "normally, Cs are Ds and we have a probability of $1 - p$ of having exceptional Cs not being Ds". Probabilities of exceptions are then used in order to reason about plausible scenarios, obtained by selecting only *some* typicality assumptions and whose probabilities belong to a given and fixed range. As a difference with the logic $\mathbf{T}^{\mathsf{CL}}$, all typicality assumptions are systematically taken into account: as a consequence, one cannot exploit such a DL for capturing compositionality, since it is not possible to block inheritance of prototypical properties in concept combination. The logic $\mathbf{T}^{\mathsf{CL}}$ extends the work of [Pozzato (2017b)] in that it does not systematically take into account all typicality assumptions. As a consequence, $\mathbf{T}^{\mathsf{CL}}$ allows to block inheritance of prototypical properties in concept combination. The same criticism applies also to the approach proposed in [Pozzato (2017a)], where $\mathcal{ALC}+\mathbf{T_R}$ is extended by inclusions of the form $\mathbf{T}(C) \sqsubseteq_d D$, where d is a *degree of expectedness*, used to define a preference relation among extended ABoxes: entailment of queries is then restricted to ABoxes that are minimal with respect to such preference relations and that represent surprising scenarios. Also in this case, however, the resulting logic does not allow to define scenarios containing only some inclusions, since all of them are systematically considered. Similarly, probabilistic DLs [Riguzzi et al. (2015a)] themselves cannot be employed as a framework for dealing with the combination of prototypical concepts, since these logics are not equipped with reasoning mechanisms needed for typicality based reasoning (and that, in $\mathbf{T}^{\mathsf{CL}}$, are inherited by $\mathcal{ALC} + \mathbf{T_R}$).

Several approaches have been recently introduced in order to tackle the problem of reasoning under probabilistic uncertainty in Description Logics. In [Lukasiewicz and Straccia (2009)] the authors combine fuzzy Description Logics, fuzzy logic programs and probabilistic uncertainty. An extension of the lightweight Description Logic $DL - Lite$ within a possibility theory setting is proposed in [Benferhat and Bouraoui (2017)]. In this approach, uncertainty degrees are associated to inclusions in order to define the inconsistency degree of a KB. All these approaches neglect to consider the proposed frameworks as the basis for the combination of concepts since, as already mentioned, approaches based on fuzziness fail to do the job.

For what concerns, more specifically, the modelling of prototypical concept composition in a human-like fashion (and with human-level performances), several approaches have been proposed in both the AI and computational cognitive science communities. Lewis and Lawry [Lewis and Lawry (2016)] present a detailed analysis of the limits of the set-theoretic approaches [Montague (1973)], the fuzzy logics [Dubois and Prade (1997); Zadeh (1975)] (whose limitations was already shown in [Hampton (2011); Osherson and Smith (1981); Smith and Osherson (1984)]), the vector-space models [Mitchell and Lapata (2010)] and quantum probability approaches [Aerts et al. (2013)] proposed to model this phenomenon. In addition, they propose to use hierarchical conceptual spaces [Gärdenfors (2014)] to model the phenomenon in a way that accurately reflects how humans exploit their creativity in conjunctive concept combination. While we agree with the authors with the comments moved

to the described approaches, in this work we have shown that our logic can equally model, in a cognitively compliant-way, the composition of prototypes by using a computationally effective nonmonotonic formalism. In particular, our model is able to meet all following cognitive requirements [Hampton (2011); Lewis and Lawry (2016)]: i) it provides a blocking mechanism of property inheritance for prototypical concept combination thus enabling the possibility of dealing with a non-standard compositional behavior ii) it is able to deal with the phenomenon of attribute emergence (and loss) for the combined concept iii) it preservers the notions of necessity and impossibility of property attribution for the combined concept iv) it explicitly assumes that the combination is not commutative (i.e. the different attribution of the HEAD-MODIFER roles does non provide the same combined concept) and that v) there are dominance effects in the concepts to be combined (both these effects are obtained via the HEAD-MODIFIER heuristics).

With respect to other formal approaches developed to model the same phenomenon [Kamp and Partee (1995); Smith et al. (1988)], we have also shown that our formalism is able to account for forms of NOUN-NOUN concept combination (e.g. like stone lion or, similarly, like porcelain cat) that such frameworks are not able to model (see [Gärdenfors (1998)]). Also: our framework does not provide any increase in the reasoning complexity with respect to the standard monotonic DL $\mathcal{ALC}$.

It is also worth-noticing that the overall reasoning procedure of $\mathbf{T}^{\mathsf{CL}}$ is directly inspired to the Composite Prototype Model (CPM) proposed by Hampton in order to account for the phenomenon of typicality-based composition [Hampton (1987, 1988, 2017)]. The similarity to CPM concerns both the identification of dominant conceptual features to consider for the composition (and this is reflected, in $\mathbf{T}^{\mathsf{CL}}$, by considering: i) the distinction between typical and rigid properties ii) the probability values associated to every typical inclusions iii) the HEAD-MODIFIER difference) and the so called phenomenon of *attribute emergence* (shown in the *Pet Fish* case). As a difference with respect to the Hampton's model, the reasoning procedure of $\mathbf{T}^{\mathsf{CL}}$ does not assume a process proceeding, first, on the attempt of obtaining the compound by composing the features and then via a process of successive identification and amalgamation of the incompatibility eventually occurring. In our case, in fact, the logic $\mathbf{T}^{\mathsf{CL}}$ obtains the combination by directly excluding the inconsistent scenarios and resorts to all its constituent ingredients (i.e. the non monotonic procedure of $\mathcal{ALC} + \mathbf{T_R}$, the probabilistic ranking provided by the DISPONTE semantics and the HEAD-MODIFIER heuristics) to determine which features to inherit.

Other attempts similar to the one proposed here concerns the modelling of the conceptual blending phenomenon: a task where the obtained concept is *entirely novel* and has no strong association with the two base concepts (for details about the differences between conceptual combination and conceptual blending see [Nagai and Taura (2006)]). In this setting, [Confalonieri et al. (2016)] proposed a mechanism for conceptual blending based on the DL $\mathcal{EL}^{++}$. They construct the generic space of two concepts by introducing an upward refinement operator that is used for finding common generalizations of $\mathcal{EL}^{++}$ concepts. However, differently from us, what they call prototypes are expressed in the standard monotonic DL, which does not allow to reason about typicality and defeasible inheritance. More recently, a different approach is proposed in [Eppe et al. (2018)], where the authors see the problem of concept blending as a nonmonotonic search problem and proposed to use Answer Set Programming (ASP) to deal with this search problem. As we have shown in the Chimera example, $\mathbf{T}^{\mathsf{CL}}$ is flexible enough to be applied also to the case of conceptual blending. There is no evidence, however, that both the frameworks of [Confalonieri et al. (2016)] and [Eppe et al. (2018)] would be able to model (in toto or in part) conceptual combination problems like the PET FISH. As such, $\mathbf{T}^{\mathsf{CL}}$ seems to provide a more general mechanism for modelling the combinatorial phenomenon of concept invention (that can be obtained both with combination and

blending).

We are currently developing an efficient reasoner for the logic $\mathbf{T}^{\mathsf{CL}}$, relying on the prover RAT-OWL [Giordano et al. (2017)] for reasoning in the nonmonotonic logic $\mathcal{ALC} + \mathbf{T_R}$ underlying our approach and on the well established HermiT reasoner. The first version of the system is implemented in Pyhton and exploits a translation of an $\mathcal{ALC} + \mathbf{T_R}$ knowledge base into standard $\mathcal{ALC}$.

In future research we aim at extending our approach to more expressive DLs, such as those underlying the standard OWL language. Starting from the work of [Giordano et al. (2014)], applying the logic with the typicality operator and the rational closure to $\mathcal{SHIQ}$, we intend to study whether and how $\mathbf{T}^{\mathsf{CL}}$ could provide an alternative solution to the problem of the "all or nothing" behavior of rational closure with respect to property inheritance.

We envision different areas of application for our framework. We have already mentioned the field of computational creativity, but other employments can be considered. For example, we plan to use the logic $\mathbf{T}^{\mathsf{CL}}$ as the basis for the exploitation of a new area of autonomic computing [Huebscher and McCann (2008)] concerning the problem of the automatic generation of novel knowledge in a cognitive artificial agent, starting from an initial commonsense knowledge base. This approach will require to enrich the knowledge processing mechanisms of current cognitive agents [Lieto et al. (2018)] and presents an element of innovation in that it does not assume that the only way to process and reason on new knowledge is via an external injection of new information but via a process of automatic knowledge generation. We believe that such an approach can find practical applications in the areas of cognitive agent architectures and robotics. Finally, as a short-term future work, we aim at investigating the use of $\mathbf{T}^{\mathsf{CL}}$ for generating metaphors based on frame-semantics (and therefore starting from event-centric representations instead of subject-centric ones) by using the recently developed repository Amnestic Forgery [Gangemi et al. (2018)]. Should $\mathbf{T}^{\mathsf{CL}}$ prove to be able to model also metaphorical phenomena starting from a different semantic approach, this would represent an additional symptom that its underlying procedure captures some foundational elements of commonsense concept combination and invention.

References

Aerts, D., Gabora, L., and Sozzo, S. (2013). Concepts and their dynamics: A quantum-theoretic modeling of human thought. *Topics in Cognitive Science*, 5(4):737–772.

Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., and Nardi, D. (2003). *The description logic handbook: Theory, implementation and applications*. Cambridge university press.

Baader, F. and Hollunder, B. (1995a). Embedding defaults into terminological knowledge representation formalisms. *Journal of Automated Reasoning (JAR)*, 14(1):149–180.

Baader, F. and Hollunder, B. (1995b). Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *Journal of Automated Reasoning (JAR)*, 15(1):41–68.

Benferhat, S. and Bouraoui, Z. (2017). Min-based possibilistic dl-lite. *J. Log. Comput.*, 27(1):261–297.

Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103(1-2):347–356.

Bonatti, P. A., Faella, M., Petrova, I., and Sauro, L. (2015). A new semantics for overriding in description logics. *Artificial Intelligence*, 222:1–48.

Bonatti, P. A., Lutz, C., and Wolter, F. (2009). The Complexity of Circumscription in DLs. *Journal of Artificial Intelligence Research (JAIR)*, 35:717–773.

Casini, G. and Straccia, U. (2010). Rational closure for defeasible description logics. In Janhunen, T. and Niemelä, I., editors, *Logics in Artificial Intelligence - Proceedings of the12th European Conference (JELIA 2010)*, volume 6341 of *Lecture Notes in Computer Science (LNCS)*, pages 77–90. Springer.

Casini, G. and Straccia, U. (2013). Defeasible Inheritance-Based Description Logics. *Journal of Artificial Intelligence Research (JAIR)*, 48:415–473.

Colton, S., Wiggins, G. A., et al. (2012). Computational creativity: The final frontier? In *Ecai*, volume 12, pages 21–26. Montpelier.

Confalonieri, R., Schorlemmer, M., Kutz, O., Peñaloza, R., Plaza, E., and Eppe, M. (2016). Conceptual blending in EL++. In Lenzerini, M. and Peñaloza, R., editors, *Proceedings of the 29th International Workshop on Description Logics, Cape Town, South Africa, April 22-25, 2016.*, volume 1577 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Donini, F. M., Nardi, D., and Rosati, R. (2002a). Description logics of minimal knowledge and negation as failure. *ACM Transactions on Computational Logics (ToCL)*, 3(2):177–225.

Donini, F. M., Nardi, D., and Rosati, R. (2002b). Description logics of minimal knowledge and negation as failure. *ACM Transactions on Computational Logics (ToCL)*, 3(2):177–225.

Dubois, D. and Prade, H. (1997). The three semantics of fuzzy sets. *Fuzzy sets and systems*, 90(2):141–150.

Eppe, M., Maclean, E., Confalonieri, R., Kutz, O., Schorlemmer, M., Plaza, E., and Kühnberger, K.-U. (2018). A computational framework for conceptual blending. *Artificial Intelligence*, 256:105–129.

Fodor, J. A. (1981). The present status of the innateness controversy.

Franks, B. (1995). Sense generation: A "quasi-classical" approach to concepts and concept combination. *Cognitive science*, 19(4):441–505.

Frixione, M. and Lieto, A. (2011). Representing and reasoning on typicality in formal ontologies. In *Proceedings of the 7th International Conference on Semantic Systems*, pages 119–125. ACM.

Gangemi, A., Alam, M., and Presutti, V. (2018). Amnestic forgery: An ontology of conceptual metaphors. In *Formal Ontology in Information Systems - Proceedings of the 10th International Conference, FOIS 2018, Cape Town, South Africa, 19-21 September 2018*, pages 159–172.

Gärdenfors, P. (1998). *Concept combination: a geometrical model*. Lund University.

Gärdenfors, P. (2004). *Conceptual spaces: The geometry of thought*. MIT press.

Gärdenfors, P. (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press.

Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2009). ALC+T: a preferential extension of description logics. *Fundamenta Informaticae*, 96:341–372.

Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2013a). A NonMonotonic Description Logic for Reasoning About Typicality. *Artificial Intelligence*, 195:165 – 202.

Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2013b). Minimal model semantics and rational closure in description logics. In Eiter, T., Glimm, B., Kazakov, Y., and Krötzsch, M., editors, *Informal Proceedings of the 26th International Workshop on Description Logics, Ulm, Germany, July 23 - 26, 2013*, volume 1014 of *CEUR Workshop Proceedings*, pages 168–180. CEUR-WS.org.

Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2014). Rational closure in $\mathcal{SHIQ}$. In *DL 2014, 27th International Workshop on Description Logics*, volume 1193 of *CEUR Workshop Proceedings*, pages 543–555. CEUR-WS.org.

Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2015). Semantic characterization of Rational Closure: from Propositional Logic to Description Logics. *Artificial Intelligence*, 226:1–33.

Giordano, L., Gliozzi, V., Pozzato, G. L., and Renzulli, R. (2017). An efficient reasoner for description logics of typicality and rational closure. In Artale, A., Glimm, B., and Kontchakov, R., editors, *Proceedings of the 30th International Workshop on Description Logics, Montpellier, France, July 18-21, 2017*, volume 1879 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Gleitman, L., Connolly, A. C., and Armstrong, S. L. (2012). Can prototype representations support composition and decomposition. *Oxford handbook of compositionality*, pages 418–436.

Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15(1):55–71.

Hampton, J. A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1):12.

Hampton, J. A. (2000). Concepts and prototypes. *Mind & Language*, 15(2-3):299–307.

Hampton, J. A. (2011). Conceptual combinations and fuzzy logic. *Concepts and fuzzy logic*, 209.

Hampton, J. A. (2017). Compositionality and concepts. In *Compositionality and concepts in linguistics*

and psychology, pages 95–121. Springer, Cham.

Huebscher, M. C. and McCann, J. A. (2008). A survey of autonomic computing—degrees, models, and applications. *ACM Computing Surveys (CSUR)*, 40(3):7.

Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57(2):129–191.

Lawry, J. and Tang, Y. (2009). Uncertainty modelling for vague concepts: A prototype theory approach. *Artificial Intelligence*, 173(18):1539–1558.

Lehmann, D. and Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60.

Lehmann, D. J. (1995). Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82.

Lewis, M. and Lawry, J. (2016). Hierarchical conceptual spaces for concept combination. *Artificial Intelligence*, 237:204–227.

Lieto, A., Lebiere, C., and Oltramari, A. (2018). The knowledge level in cognitive architectures: Current limitations and possible developments. *Cognitive Systems Research*, 48:39–55.

Lieto, A. and Pozzato, G. L. (2018). A description logic of typicality for conceptual combination. In *Proceedings of the 24th International Symposium on Methodologies for Intelligent Systems ISMIS2018*.

Lieto, A. and Pozzato, G. L. (2019). Applying a description logic of typicality as a generative tool for concept combination in computational creativity. *Intelligenza Artificiale*, (to appear).

Lieto, A., Radicioni, D. P., and Rho, V. (2015). A common-sense conceptual categorization system integrating heterogeneous proxytypes and the dual process of reasoning. In *In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, AAAI Press*, pages 875–881.

Lieto, A., Radicioni, D. P., and Rho, V. (2017). Dual PECCS: a cognitive system for conceptual representation and categorization. *Journal of Experimental & Theoretical Artificial Intelligence*, 29(2):433–452.

Lukasiewicz, T. (2008). Expressive probabilistic description logics. *Artificial Intelligence*, 172(6-7):852–883.

Lukasiewicz, T. and Straccia, U. (2009). Description logic programs under probabilistic uncertainty and fuzzy vagueness. *International Journal of Approximate Reasoning*, 50(6):837–853.

Mitchell, J. and Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive science*, 34(8):1388–1429.

Montague, R. (1973). The proper treatment of quantification in ordinary english. In *Approaches to natural language*, pages 221–242. Springer.

Murphy, G. (2002). *The big book of concepts*. MIT press.

Nagai, Y. and Taura, T. (2006). Formal description of concept-synthesizing process for creative design. *Design computing and cognition*, pages 443–460.

Osherson, D. N. and Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9(1):35–58.

Pelletier, F. J. (2017). *Compositionality and Concepts—A Perspective from Formal Semantics and Philosophy of Language*, pages 31–94. Springer International Publishing, Cham.

Pozzato, G. L. (2017a). Reasoning about plausible scenarios in description logics of typicality. *Intelligenza Artificiale*, 11(1):25–45.

Pozzato, G. L. (2017b). Reasoning in description logics with typicalities and probabilities of exceptions. In Antonucci, A., Cholvy, L., and Papini, O., editors, *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*, pages 409–420. Springer.

Prinz, J. (2012). Regaining composure: A defense of prototype compositionality. *The Oxford handbook of compositionality*, pages 437–453.

Riguzzi, F., Bellodi, E., Lamma, E., and Zese, R. (2015a). Probabilistic description logics under the distribution semantics. *Semantic Web*, 6(5):477–501.

Riguzzi, F., Bellodi, E., Lamma, E., and Zese, R. (2015b). Reasoning with probabilistic ontologies. In Yang, Q. and Wooldridge, M., editors, *Proceedings of the Twenty-Fourth International Joint*

Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015, pages 4310–4316. AAAI Press.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of experimental psychology: General*, 104(3):192.

Sato, T. (1995). A statistical learning method for logic programs with distribution semantics. In Sterling, L., editor, *Logic Programming, Proc. of the 12th International Conference on Logic Programming*, pages 715–729. MIT Press.

Smith, E. E. and Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive science*, 8(4):337–361.

Smith, E. E., Osherson, D. N., Rips, L. J., and Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive science*, 12(4):485–527.

Tversky, A. and Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review*, 90(4):293.

Zadeh, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning—i. *Information sciences*, 8(3):199–249.