COCOS: a typicality based COncept COmbination System

Antonio Lieto, Gian Luca Pozzato, and Alberto Valese

Dipartimento di Informatica, Università di Torino, Italy
{antonio.lieto@,gianluca.pozzato@,alberto.valese@edu}.unito.it

Abstract. In this short paper we describe COCOS, a tool we are currently developing in order to account for the phenomenon of combining prototypical concepts, an open problem in the fields of AI and cognitive modelling. COCOS is based on a probabilistic extension of the logic of typicality $\mathcal{ALC} + \mathbf{T_R}$ by inclusions $p :: \mathbf{T}(C) \sqsubseteq D$ ("we have probability p that typical Cs are Ds") and it embeds a set of cognitive heuristics for concept combination.

1 Introduction

Inventing novel concepts by combining the typical knowledge of pre-existing ones is an important human creative ability. Dealing with this problem requires, from an AI perspective, the harmonization of two conflicting requirements that are hardly accommodated in symbolic systems: the need of a syntactic compositionality (typical of logical systems) and that one concerning the exhibition of typicality effects [1]. According to a well-known argument [2], in fact, prototypical concepts are not compositional. The argument runs as follows: consider a concept like *pet fish*. It results from the composition of 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 paper we describe our work in progress in this field of research. We describe COCOS, a software system able to account for this type of human-like concept combination. COCOS relies on a nonmonotonic Description Logic (DL) of typicality called $\mathbf{T}^{\mathsf{CL}}$ (typical compositional logic) recently introduced in [3]. This logic, whose complexity is ExpTime-complete as the underlying standard $\mathcal{ALC}$, combines two main ingredients. The first one relies on the DL of typicality $\mathcal{ALC} + \mathbf{T_R}$ introduced in [4]. 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 $\mathbf{T}(Athlete) \sqsubseteq InFit$ and $\mathbf{T}(SumoWrestler) \sqsubseteq \neg InFit$, given that $SumoWrestler \sqsubseteq Athlete$. The semantics of the $\mathbf{T}$ operator is characterized by the properties of *rational logic* [5], 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 [5] 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 $\mathbf{T}(Athlete \sqcap Bald) \sqsubseteq InFit$ (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 consider a distributed semantics similar to the one of probabilistic DLs known as DISPONTE [6], allowing one to label axioms with degrees representing probabilities, but restricted to typicality inclusions. The basic idea is to label inclusions $\mathbf{T}(C) \sqsubseteq D$ with a real number between 0.5 and 1, representing its probability, assuming that each axiom is independent from each others. 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. In a slight extension of the above example, we could have the need of representing that both the typicality 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) $0.8 \ :: \ \mathbf{T}(Athlete) \sqsubseteq InFit$; (3) $0.8 \ :: \ \mathbf{T}(SumoWrestler) \sqsubseteq \neg InFit$; (4) $0.95 \ :: \ \mathbf{T}(Athlete) \sqsubseteq YoungPerson$. We consider eight different scenarios, representing all possible combinations of typicality inclusion: as an example, $\{((2), 1), ((3), 0), ((4), 1)\}$ represents the scenario in which (2) and (4) hold, whereas (3) does not. We equip each scenario with a probability depending on those of the involved inclusions, then we restrict reasoning to scenarios whose probabilities belong to a given and fixed range.

The proposed system COCOS is able to tackle the problem of composing prototypical concepts. As an additional element we employ a method inspired by cognitive semantics [7] for the identification of a dominance effect between the concepts to be combined: for every combination, we distinguish a HEAD, representing the stronger element of the combination, and a MODIFIER. The basic idea is: given a KB and two concepts C_H (HEAD) and C_M (MODIFIER) occurring in it, we consider only *some* scenarios in order to define a revised knowledge base, enriched by typical properties of the combined concept $C \sqsubseteq C_H \sqcap C_M$. We use COCOS for the generation of novel creative concepts, that could be useful in many applicative scenarios.

2 The COCOS system

The current version of the system is implemented in Pyhton and exploits the translation of an $\mathcal{ALC} + \mathbf{T_R}$ knowledge base into standard $\mathcal{ALC}$ introduced in [4] and adopted by the system RAT-OWL [8]. COCOS makes use of the library `owlready2` [1] that allows one to rely on the services of efficient DL reasoners, e.g. the HermiT reasoner.

Our system relies on the procedures developed for logic $\mathbf{T}^{\text{CL}}$. More in detail, we consider a KB $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ 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

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

[1] https://pythonhosted.org/Owlready2/

where $p \in (0.5, 1) \subseteq \mathbb{R}$ is the probability of the typicality inclusion and $\mathcal{A}$ is the ABox, i.e. a finite set of formulas of the form either $C(a)$ or $R(a, b)$, where a and b are individual names, C is a concept and R is a role.

Given two concepts C_H and C_M occurring in $\mathcal{K}$, the logic $\mathbf{T}^{\textsf{CL}}$ allows one to define the compound concept C as the combination of the HEAD C_H and the MODIFIER C_M, where $C \sqsubseteq C_H \sqcap C_M$ and the typical properties of the form $\mathbf{T}(C) \sqsubseteq D$ to ascribe to the concept C are obtained by selecting suitable *scenarios*. 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 coming from either the HEAD or the MODIFIER. Each scenario is equipped by a probability defined as $p_1 \times \ldots \times p_i \times (1 - q_1) \times \ldots \times (1 - q_j)$, where $p_1 :: \mathbf{T}(E_1) \sqsubseteq F_1, \ldots, p_i :: \mathbf{T}(E_i) \sqsubseteq F_i$ are the typicalities included in the scenario, whereas $q_1 :: \mathbf{T}(G_1) \sqsubseteq H_1, \ldots, q_j :: \mathbf{T}(G_j) \sqsubseteq H_j$ are those *not* included in the scenario.

Intuitively, the selected scenarios are those satisfying the following properties:

1. are consistent with respect to $\mathcal{K}$;
2. are not trivial, in the sense that the scenarios considering *all* properties that can be consistently ascribed to C are discarded;
3. are those giving preference to the typical properties of the HEAD C_H (with respect to those of the MODIFIER C_M). Notice that, in case of conflicting properties like D and $\neg D$, given two scenarios w_1 and w_2 such that an inclusion $p_1 :: \mathbf{T}(C_H) \sqsubseteq D$ belongs to w_1 whereas $p_2 :: \mathbf{T}(C_M) \sqsubseteq \neg D$ belongs to w_2, the scenario w_2 is discarded in favor of w_1.

In order to select the wanted scenarios, COCOS applies 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, COCOS first discards all the inconsistent scenarios, then it considers the remaining (consistent) ones in decreasing order by their probabilities. It then considers the blocks of scenarios with the same probability as follows:

- it discards scenarios considered as *trivial*, consistently inheriting all (or most of) the properties from the starting concepts to be combined;
- among the remaining ones, it discards those inheriting properties from the MODIFIER in conflict with properties inherited from the HEAD in another scenario of the same block (i.e., with the same probability);
- 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, COCOS repeats the procedure by considering the block of scenarios, all having the immediately lower probability.

The output of COCOS corresponds to the set of remaining scenarios.

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:

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

for all D such that $\mathbf{T}(C) \sqsubseteq D$ is entailed in all the scenarios selected by COCOS. Notice that, since the *C-revised* knowledge base is still in the language of the $\mathbf{T}^{\textsf{CL}}$

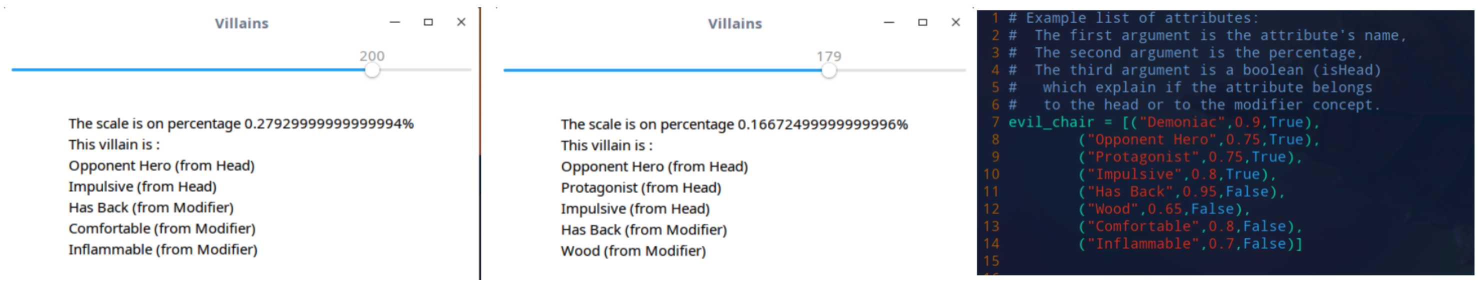

Fig. 1. Some pictures of COCOS. In addition to presenting the selected scenario with typical properties of the combined concept (right on the figure), a slider allows the user to select alternative scenarios, ranging from more trivial to more surprising ones (center and left on the figure).

logic, we can iteratively repeat the same procedure in order to combine not only atomic concepts, but also compound concepts.

Some pictures of COCOS are shown in Figure 2. An example of the application of COCOS for the generation of a new concept is shown in the next concluding section.

3 An Example of Application of COCOS

We exploit the system COCOS as a creative support tool to generate a new type of character in the field of computational creativity for a video game or a movie.

Let us assume to generate a novel concept obtained as the combination of concepts $Villain$ (as HEAD) and $Chair$ (as MODIFIER). Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \emptyset \rangle$ be as follows:

R1 $Villain \sqsubseteq \exists fightsFor.PersonalGoal$
R2 $Villain \sqsubseteq Animate$
R3 $Villain \sqsubseteq \exists hasValues.NegativeMoralValues$
R4 $Chair \sqsubseteq hasComponent.SupportingSitComponent$
R5 $Chair \sqsubseteq hasComponent.Sit$
R6 $CollectiveGoal \sqcap PersonalGoal \sqsubseteq \bot$

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

T1 $0.9 \ :: \ \mathbf{T}(Villain) \sqsubseteq \exists hasIconicity.Demoniac$
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 $0.95 \ :: \ \mathbf{T}(Chair) \sqsubseteq hasComponent.Back$
T7 $0.65 \ :: \ \mathbf{T}(Chair) \sqsubseteq madeOf.Wood$
T8 $0.8 \ :: \ \mathbf{T}(Chair) \sqsubseteq Comfortable$
T9 $0.7 \ :: \ \mathbf{T}(Chair) \sqsubseteq Inflammable$

We consider the 512 scenarios that can be generated from nine typicality inclusions T1–T9, 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 $\mathit{Villain} \sqcap \mathit{Chair} \sqsubseteq \mathit{Animate}$, therefore all scenarios including T5, imposing that $\mathit{Villain} \sqcap \mathit{Chair} \sqsubseteq \neg \mathit{Animate}$ are inconsistent. We also discard the most obvious scenario including all the typicality inclusions of $\mathcal{T}$, having probability of 14%. Among the remaining scenarios, the most probable contains all the inclusions related to the HEAD, namely T1, T2, T3, and T4, whereas it contains T6, T8, and T9 concerning the MODIFIER. This scenario is the preferred one from a cognitive point of view. 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, by imposing that the compound concept should inherit six properties, 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, Similarly, we could decide to prefer - still more surprising - scenarios, by selecting those with probability of 2.51%, obtaining the following, plausible but not obvious, creative definitions of villain chair:

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

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

References

1. Frixione, M., Lieto, A.: Representing concepts in formal ontologies: Compositionality vs. typicality effects. Logic and Logical Philosophy **21**(4) (2012) 391–414
2. Osherson, D.N., Smith, E.E.: On the adequacy of prototype theory as a theory of concepts. Cognition **9**(1) (1981) 35–58
3. Lieto, A., Pozzato, G.L.: A description logic of typicality for conceptual combination. In: Proceedings of the 24th International Symposium on Methodologies for Intelligent Systems ISMIS2018. (2018)
4. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Semantic characterization of Rational Closure: from Propositional Logic to Description Logics. Artif. Intelligence **226** (2015) 1–33
5. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**(1) (1992) 1–60
6. Riguzzi, F., Bellodi, E., Lamma, E., Zese, R.: Reasoning with probabilistic ontologies. In Yang, Q., Wooldridge, M., eds.: Proceedings of IJCAI 2015. 4310–4316
7. Hampton, J.A.: Inheritance of attributes in natural concept conjunctions. Memory & Cognition **15**(1) (1987) 55–71
8. Giordano, L., Gliozzi, V., Pozzato, G.L., Renzulli, R.: An efficient reasoner for description logics of typicality and rational closure. In: Proc. of DL 2017. Volume 1879 of CEUR. (2017)