AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Beyond subgoalings: A dynamic knowledge generation framework for creative problem solving in cognitive architectures

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1726157> since 2020-01-31T15:55:35Z

Published version:

DOI:10.1016/j.cogsys.2019.08.005

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)

Beyond Subgoalng: A Dynamic Knowledge Generation Framework for Creative Problem Solving in Cognitive Architectures

Antonio Lieto^{a,b}, Federico Perrone^a, Gian Luca Pozzato^a, Eleonora Chiodino^c

^aUniversity of Turin, Dept. of Computer Science, Turin, Italy

^bICAR-CNR, Palermo, Italy

^cUniversity of Turin, Dept. of Psychology, Turin, Italy

P.S. This is a pre-print version. The final version is available in the Cognitive Systems Research journal at <https://doi.org/10.1016/j.cogsys.2019.08.005>

Abstract

In this paper we propose a computational framework aimed at extending the problem solving capabilities of cognitive artificial agents through the introduction of a novel, goal-directed, dynamic knowledge generation mechanism obtained via a non monotonic reasoning procedure. In particular, the proposed framework relies on the assumption that certain classes of problems cannot be solved by simply learning or injecting new external knowledge in the declarative memory of a cognitive artificial agent but, on the other hand, require a mechanism for the automatic and creative re-framing, or re-formulation, of the available knowledge. We show how such mechanism can be obtained through a framework of dynamic knowledge generation that is able to tackle the problem of commonsense concept combination. In addition, we show how such a framework can be employed in the field of cognitive architectures in order to overcome situations like the *impasse* in SOAR by extending the possible options of its subgoalng procedures.

Keywords:

Knowledge Representation, Cognitive Architectures, Knowledge Generation, Concept Combination, Commonsense reasoning

1. Introduction

Goal-directed problem solving is a crucial everyday activity for both natural and artificial systems. A straightforward assumption in goal-directed systems is that, in the

cases where a given goal cannot be reached, a replanning strategy is required in order to change the original goal and/or reconfigure the set of actions originally selected to perform that goal [2]. Usually such goal reconfiguration is based on the availability of novel, additional, knowledge that can be then used to select novel sub-goals or novel operations to carry on. In this work, we consider those situations where the solution to a given problem cannot come with the classical means usually adopted for obtaining new knowledge (and leading to a goal-redefinition). In particular, we consider scenarios where the availability of novel knowledge cannot be obtained in an *extrinsic* way (e.g. via communication with another agent or, via a novel learning process or by an external injection of novel knowledge in the declarative memory of an artificial system). On the other hand, in such scenarios, the key to the problem solution lies in an *intrinsic* agent capability of automatically generating novel knowledge by recombining, in a dynamic and innovative way, the possessed knowledge in order to look with new eyes to the problem in hand and solve it.

In this paper we present a framework for the dynamic and automatic generation of novel knowledge obtained through a process of commonsense reasoning based on typicality-based concept combination. We exploit a recently introduced extension of a Description Logic of typicality able to combine prototypical descriptions of concepts in order to generate new prototypical concepts. Intuitively, in the context of our application of this logic, the overall pipeline works as follows: given a goal expressed as a set of properties, if the knowledge base does not contain a concept able to fulfill all these properties, then our system looks for two concepts to recombine in order to extend the original knowledge base and satisfy the goal.

The rest of the paper is organized as follows. In section 2 we describe the rationale of our proposal. In section 3 we formally describe the Description Logic of typicality used for generating novel knowledge in order to achieve a given goal by combining commonsense representations of concepts. In section 4 we describe the system adopting the proposed logic, whose efficacy is tested in section 5 in the task of object composition. In section 6 we show how the proposed framework is compliant with all the major SOAR mechanism and can be used to extend its subgoaling procedures. Finally, in section 7 we survey related approaches and conclude with a discussion on future

works.

2. Commonsense Concept Invention via Dynamic Knowledge Combination

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 of artificial intelligence [6]. Dealing with this problem requires, from an AI and cognitive modelling perspective, the harmonization of two conflicting requirements that are hardly accommodated in symbolic systems [10]): 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 [42], in fact, prototypes (i.e. commonsense conceptual representations based on typical properties) 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 work we exploit a framework able to account for this type of human-like concept combination and propose to use it as a novel mechanism able to expand the spectrum of subgoaling procedures in cognitive artificial systems. In particular, we adopt a nonmonotonic extension of Description Logics (from now on DL) ¹ able to reason on typicality and called $\mathbf{T}^{\text{CL}}$ (typicality-based

¹Description Logics are a class of decidable fragments of first order logics that are at the base of Ontology Web Language (OWL) used for the realization of computational ontologies. Nowadays DLs are the most important and widespread symbolic knowledge-representation systems. Their success is justified, on the one hand, by the fact that DLs have a well defined semantics and, on the other hand, by the fact that they offer a good trade-off between expressivity of the language and computational complexity. According to DLs, a knowledge base contains two components: 1. a TBox, containing inclusion relations among concepts, for instance $Cat \sqsubseteq Mammal$ (cats are mammals); 2. an ABox, containing facts about individuals, for instance $Cat(tom)$ (Tom is a cat). We remind to [4] for a complete introduction to DLs.

compositional logic) introduced in [29].

This logic combines two main ingredients. The first one relies on the DL of typicality $\mathcal{ALC} + \mathbf{T}_R$ introduced in [16], which allows to describe the *prototype* of a concept. 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 C s are also D s”. 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 fit”, whereas “sumo wrestlers usually are not fit” by $\mathbf{T}(\textit{Athlete}) \sqsubseteq \textit{Fit}$ and $\mathbf{T}(\textit{SumoWrestler}) \sqsubseteq \neg \textit{Fit}$, given that $\textit{SumoWrestler} \sqsubseteq \textit{Athlete}$. The semantics of the $\mathbf{T}$ operator is characterized by the properties of *rational logic* [23], 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 [23] 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 $\mathbf{T}(\textit{Athlete} \sqcap \textit{Bald}) \sqsubseteq \textit{Fit}$ (being bald is irrelevant with respect to being fit) and, if one knows that Hiroyuki is a typical sumo wrestler, to infer that he is not 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 [43], allowing to label inclusions $\mathbf{T}(C) \sqsubseteq D$ with a real number between 0.5 and 1, representing its degree of belief/probability, assuming that each axiom is independent from each others. Degrees of belief in typicality inclusions allow to define a probability distribution over *scenarios*: roughly speaking, a scenario is obtained by choosing, for each typicality inclusion, whether it is considered as true or false². In a slight extension of the above example, we could have

² We do not employ the whole characteristics of DISPONTE: as far as the inferential capabilities are

the need of representing that both the typicality inclusions about athletes and sumo wrestlers have a degree of belief of 80%, whereas we also believe that athletes are usually young with a higher degree of 95%, with the following KB: (1) $\text{SumoWrestler} \sqsubseteq \text{Athlete}$; (2) $0.8 \quad :: \quad \mathbf{T}(\text{Athlete}) \sqsubseteq \text{Fit}$; (3) $0.8 \quad :: \quad \mathbf{T}(\text{SumoWrestler}) \sqsubseteq \neg \text{Fit}$; (4) $0.95 \quad :: \quad \mathbf{T}(\text{Athlete}) \sqsubseteq \text{YoungPerson}$. 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: the scenario of the example, has probability 0.8×0.95 (since 2 and 4 are involved) $\times (1 - 0.8)$ (since 3 is not involved) $= 0.152 = 15.2\%$. Such probabilities are then taken into account in order to choose the most adequate scenario describing the prototype of the combined concept. As an additional element of the proposed formalization we employ a method inspired by cognitive semantics [18] 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 $C \sqsubseteq C_H \sqcap C_M$. In [29], we have shown that the procedure for combining concepts in the logic $\mathbf{T}^{\text{cl}}$ is essentially inexpensive, in the sense that reasoning in this logic is EXPTIME-complete as for the underlying standard Description Logic $\mathcal{ALC}$.

We extend a preliminary result by [29], and we show that the logic $\mathbf{T}^{\text{cl}}$ can be used to extend the problem solving and subgoaling procedure of cognitive agents.

3. A Non Monotonic Reasoning Framework for Concept Combination

In this section, we describe the nonmonotonic Description Logic $\mathbf{T}^{\text{cl}}$ that combines the semantics based on the rational closure of $\mathcal{ALC} + \mathbf{T}_R$ [16] with the DISPONTE semantics [43] of probabilistic DLs.

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.

By taking inspiration from [33], in our representational assumptions we consider two types of properties associated to a given concept: rigid and typical. Rigid properties are those defining a concept, e.g. $C \sqsubseteq D$ (all C s are D s). 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)³.

The language of $\mathbf{T}^{\text{cl}}$ 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, C s are also D ” (the reason why we only allow typicality inclusions equipped with $p > 0.5$ is detailed below).

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

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

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

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

³It is worth-noting that a general framework for the automatic identification of a HEAD/MODIFIER combination is currently 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 describe the reasoning part.

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 $a, b \in \mathcal{O}$ and $R \in \mathcal{R}$.

As mentioned, following the DISPONTE semantics, in $\mathbf{T}^{\text{cl}}$ each typicality inclusion is independent from each others. This avoids the problem of dealing with degrees of inconsistent inclusions.

As indicated 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. On the other hand, since the very cognitive notion of typicality derives from that one of probability distribution [44], 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 [22]), we only allow typicality inclusions equipped with degrees of belief $p > 0.5$. For such reason, in our effort of integrating two different semantics – DISPONTE and typicality logic – the choice of having degrees higher than 0.5 for typicality inclusions seems to be the only one compliant with both the formalisms. In fact, despite the DISPONTE semantics in [43] allows to assign also low degrees of belief to standard inclusions, in the logic $\mathbf{T}^{\text{cl}}$, for what explained above, it would be misleading to also allow low degrees of belief for typicality inclusions. For example, the logic $\mathbf{T}^{\text{cl}}$ does not allow an inclusion like $0.3 :: \mathbf{T}(\text{Student}) \sqsubseteq \text{YoungPerson}$, 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 $\mathbf{T}^{\text{cl}}$: 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 $0.8 :: \mathbf{T}(\text{Student}) \sqsubseteq \neg \text{Married}$, rather than $0.2 :: \mathbf{T}(\text{Student}) \sqsubseteq \text{Married}$.

A model $\mathcal{M}$ in the logic $\mathbf{T}^{\text{cl}}$ extends standard $\mathcal{ALC}$ models by a preference relation among domain elements as in the logic of typicality [16]. 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 normal

than x . Formally:

Definition 3.2 (Model of $\mathbf{T}^{\text{cl}}$). A model $\mathcal{M}$ is any structure $\langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$ where: (i) $\Delta^{\mathcal{I}}$ is a non empty set of items called the domain; (ii) $<$ 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 $\Delta^{\mathcal{I}}$; (iii) $.^{\mathcal{I}}$ 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 usual for concepts of $\mathcal{ALC}$, whereas for typicality concepts we define $(\mathbf{T}(C))^{\mathcal{I}} = \text{Min}_{<}(C^{\mathcal{I}})$, where $\text{Min}_{<}(C^{\mathcal{I}}) = \{x \in C^{\mathcal{I}} \mid \nexists y \in C^{\mathcal{I}} \text{ such that } y < x\}$.

A model $\mathcal{M}$ can be equivalently defined by postulating the existence of a function $k_{\mathcal{M}} : \Delta^{\mathcal{I}} \mapsto \mathbb{N}$, where $k_{\mathcal{M}}$ assigns a finite rank to each domain element: the rank of x is the length of the longest chain $x_0 < \dots < x$ 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 $k_{\mathcal{M}}(x) < k_{\mathcal{M}}(y)$.

Given a KB $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ and 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 $\mathcal{O}$. We say that (i) $\mathcal{M}$ satisfies $\mathcal{R}$ if, for all $C \sqsubseteq D \in \mathcal{R}$, we have $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$; (ii) $\mathcal{M}$ satisfies $\mathcal{T}$ if, for all $q :: \mathbf{T}(C) \sqsubseteq D \in \mathcal{T}$, we have that $\mathbf{T}(C)^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, i.e. $\text{Min}_{<}(C^{\mathcal{I}}) \subseteq D^{\mathcal{I}}$; (iii) $\mathcal{M}$ 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 $\text{KB} \subseteq \text{KB}'$, i.e. the resulting logic is monotonic. In order to perform useful nonmonotonic inferences, in [16] the authors have strengthened the semantics of the DL with typicality by restricting entailment to a class of minimal models. Intuitively, the idea is to restrict entailment to models that *minimize the untypical instances of a concept*. The resulting logic $\mathcal{ALC} + \mathbf{T}_R$ corresponds to a notion of *rational closure* for DLs as a natural extension of the construction provided in [23] 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 the individuals that enjoy properties that are consistent with the 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 [16] it is shown that query entailment in the nonmonotonic $\mathcal{ALC} + \mathbf{T}_R$ is in EXPTIME.

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 a subset of typicality inclusions in $\mathcal{T}$.

Definition 3.3 (Atomic choice). Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 :: \mathbf{T}(C_1) \sqsubseteq D_1, \dots, E_n = q_n :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ we define (E_i, k_i) an atomic choice, where $k_i \in \{0, 1\}$.

Definition 3.4 (Selection). Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 :: \mathbf{T}(C_1) \sqsubseteq D_1, \dots, E_n = q_n :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ 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 $i = 1, 2, \dots, n$. 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 3.5 (Scenario). Given $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T} = \{E_1 = q_1 :: \mathbf{T}(C_1) \sqsubseteq D_1, \dots, E_n = q_n :: \mathbf{T}(C_n) \sqsubseteq D_n\}$ 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 $\mathbf{T}^{\text{cl}}$ satisfying $\mathcal{K}$.

We denote with $\mathcal{W}_{\mathcal{K}}$ the set of all scenarios. 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 $\mathcal{K}$, the logic $\mathbf{T}^{\text{CL}}$ allows defining a prototype of 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 by considering blocks of scenarios with the same probability, in decreasing order starting from the highest one. We first discard all the inconsistent scenarios, then:

- we discard those scenarios considered as *trivial*, consistently inheriting all the properties from the HEAD from the starting concepts to be combined. This choice is motivated by the challenges provided by task of commonsense conceptual combination itself: in order to generate plausible and creative compounds it is necessary to maintain a level of surprise in the combination. Thus both scenarios inheriting all the properties of the two concepts and all the properties of the HEAD are discarded since prevent this surprise;
- 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, having the immediately lower probability.

Remaining scenarios with the highest probability are those selected by the logic $\mathbf{T}^{\text{CL}}$. The ultimate output of our mechanism is a knowledge base in the logic $\mathbf{T}^{\text{CL}}$ 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 from w in the logic $\mathbf{T}^{\text{CL}}$. 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 degree of belief 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 degree of belief 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 , and p is defined as above.

In [29] we have shown that reasoning in $\mathbf{T}^{\text{CL}}$ remains in the same complexity class of standard $\mathcal{ALC}$ Description Logics.

Theorem 3.6. *Reasoning in $\mathbf{T}^{\text{CL}}$ is EXPTIME-complete.*

4. A Goal-directed System for Dynamic Knowledge Generation and Invention

In this section we describe a goal-directed system relying on the above illustrated $\mathbf{T}^{\text{CL}}$ logic ⁴. In particular, the system (available at the URL: <http://di.unito.it/GOCCIOLA>) is able to dynamically generate novel knowledge in the cases in which the original goal cannot be directly solved by a given agent only by resorting to its available knowledge. The process of automatic knowledge generation, as mentioned, is obtained by adopting the process of commonsense concept combination of $\mathbf{T}^{\text{CL}}$, namely: by combining concepts in the knowledge base which are relevant for the task to solve.

The overall pipeline of the system can be described as follows: the system receives in input a certain goal to achieve. The goal is expressed in terms of tuples representing

⁴In other works we have already shown how such logic can be used to model complex cognitive phenomena [31] (including metaphors generation) and to build intelligent applications in the field of computational creativity [30].

the desired final state. For example: a goal can be expressed as $\{Object, Cutting, Graspable\}$ to identify the scope of retrieving, from the inventory of the available knowledge in the agent declarative memory, an element that is a graspable object able to cut some surfaces. Once processed the input, the system verifies, via a searching process in the hybrid, probabilistic, knowledge base assumed in $\mathbf{T}^{cl}$, whether there is some element that can directly satisfy the desired conditions. If so, the element(s) (if any) satisfying the request are returned and ranked in descending order of probability. If not, the system tries to perform a task of semantic-driven goal-reformulation by looking for WordNet synonyms and hyperonyms⁵ of the terms specified in input (in order to find at least a minimal set of candidate concepts sharing, if considered jointly, all the required goal desiderata). Once this process is also executed, and the minimal set of candidate concepts that (jointly) can be combined to satisfy the goal is reached, the system adopt the typicality-based reasoning procedure of concept combination developed in $\mathbf{T}^{cl}$.

More formally:

Definition 4.1. *Given a knowledge base $\mathcal{K}$ in the logic $\mathbf{T}^{cl}$, let $\mathcal{G}$ be a set of concepts $\{D_1, D_2, \dots, D_n\}$ called goal. We say that a concept C is a solution to the goal $\mathcal{G}$ if either:*

- *for all $D_i \in \mathcal{G}$, either $\mathcal{K} \models C \sqsubseteq D_i$ or $\mathcal{K} \models \mathbf{T}(C) \sqsubseteq D_i$ in the logic $\mathbf{T}^{cl}$*

or

- *C corresponds to the combination of, at least, two concepts C_1 and C_2 occurring in $\mathcal{K}$, i.e. $C \equiv C_1 \sqcap C_2$, and the C -revised knowledge base $\mathcal{K}_C$ provided by the logic $\mathbf{T}^{cl}$ is such that, for all $D_i \in \mathcal{G}$, either $\mathcal{K}_C \models C \sqsubseteq D_i$ or $\mathcal{K}_C \models \mathbf{T}(C) \sqsubseteq D_i$.*

⁵WordNet is a widely known lexical database for the English language [35]. Rather than organizing terms alphabetically (like ordinary dictionaries, where senses are possibly scattered) WN groups terms into synonyms sets called *synsets*, that are equipped with short definitions and usage examples. Such sets are represented as the nodes of a large semantic network, where the intervening edges represent a number of semantic relations among synset elements (such as hyponymy, hypernymy, antonymy, meronymy, holonymy).

In case the goal cannot be achieved in a direct way (i.e. there is no element in the KB satysfying the goal desiderata) the system computes a list of concepts of the initial knowledge base satisfying at least a property of the goal (using Wordnet if the initial goal formulation does not satisfy such condition). As an example, suppose to have:

$$\mathcal{G} = \{Object, Graspable, Cutting\},$$

and suppose that the following inclusions belong to the knowledge base:

$$\begin{aligned} Spoon &\sqsubseteq Graspable \\ 0.85 &:: \mathbf{T}(Spoon) \sqsubseteq \neg Cutting \\ 0.9 &:: \mathbf{T}(Vase) \sqsubseteq Graspable \\ Vase &\sqsubseteq Object \end{aligned}$$

Both *Vase* and *Spoon* are included in the list of candidate concepts to be combined (along with other concepts satysfying, for example other properties of the goal such as, for example, being able to cut some surface). As a second step, for each item in the list of candidate concepts to be combined, the system computes a rank of the concept as the sum of the probabilities of the properties also belonging to the goal, assuming a score of 1 in case of a rigid property. In the example, *Vase* is ranked as $0.9 + 1 = 1.9$, since both *Graspable* and *Object* are properties belonging to the goal: for the former we take the probability 0.9 of the typicality inclusion $\mathbf{T}(Vase) \sqsubseteq Graspable$, for the latter we provide a score of 1 since the property $Vase \sqsubseteq Object$ is rigid. Concerning the concept *Spoon*, the system computes a rank of 1: indeed, the only inclusion matching the goal is the rigid one $Spoon \sqsubseteq Graspable$. Finally, the system checks whether the concept obtained by combining the candidate concepts with the highest ranks, (e.g. C_1 and C_2 in case of only 2 concepts), is able to satisfy the initial goal. The system computes a double attempt, by considering first C_1 as the HEAD and C_2 as the MODIFIER and, in case of failure, C_2 as the HEAD and C_1 as the MODIFIER.

In order to combine the two candidate concepts C_1 and C_2 , our system exploits COCOS [28], a tool generating scenarios and choosing the selected one(s) according

to the logic $\mathbf{T}^{\text{cl}}$. COCOS makes use of the library `owlready2`⁶ that allows one to rely on the services of efficient DL reasoners, e.g. the HermiT reasoner.

5. Experimentation

In this section, we describe the experimental setup and the obtained results of our system in task of object composition of compound tools. Such ability represents a very important creative faculty found only in primates (specifically, humans and great apes) and, more recently in ravens [46]. It still represents an open challenge in the field of AI and cognitive modelling. As we will see later in detail, in fact, a major problem consists in the lack of realistic benchmarks for evaluating the performance on this task for both humans and artificial systems (this problem is also explicitly reported in [41] that represent, to the best of our knowledge, the first attempt of modelling such faculty in an artificial system). Despite the lack of such a benchmark, for our purposes we tested our system on the same proof-of-concept evaluation presented in [41]. In addition, we also provided a comparison with responses provided by human judges for the concept composition task.

5.1. Setup

Knowledge about goals, objects and entities can be represented in the system in symbolic terms. As an example, let us consider the above mentioned goal: *object, cutting, graspable*. The initial knowledge base is formalized in the language of the logic $\mathbf{T}^{\text{cl}}$ and it is stored in a suitable file. Rigid properties, holding for all individuals of a given class, are stored as pairs object-property, whereas typical properties are formalized as triples object-property-probability. We have considered an extension with probabilities of a portion of the ontology `Open Cyc` [24]⁷ referring to physical objects and tools of ordinary use in a domestic environment (e.g. a glass, a vase etc.). The considered branch of the Cyc ontology (formalized in standard Description

⁶<https://pythonhosted.org/Owlready2/>

⁷<https://github.com/asanchez75/opencyc/blob/master/opencyc-latest.owl.gz>.

Logic and, as a consequence, not able to represent and reason on typicality-based information) has been manually extended in the language of the logic $\mathbf{T}^{\text{CL}}$. Therefore the symbolic representation of the ontological objects additionally includes the following typical and functional characteristics: color, size, function, physical affordance, shape, material. Please note that it was not mandatory to fill every property of the schema for the description of objects.

As an example, the concept *Vase* is represented as follows (on the right the corresponding knowledge base in $\mathbf{T}^{\text{CL}}$):

vase, object	$Vase \sqsubseteq Object$
vase, high convexity	$Vase \sqsubseteq HighConvexity$
vase, ceramic, 0.8	$0.8 :: \mathbf{T}(Vase) \sqsubseteq Ceramic$
vase, to put plants, 0.9	$0.9 :: \mathbf{T}(Vase) \sqsubseteq ToPutPlants$
vase, to contain objects, 0.9	$0.9 :: \mathbf{T}(Vase) \sqsubseteq ToContainObjects$
vase, graspable, 0.9	$0.9 :: \mathbf{T}(Vase) \sqsubseteq Graspable$

5.2. Results of Knowledge Generation via Concept Composition

We tested the proposed framework in the task of object composition. In particular, for this task we used the same setup adopted in [41] by using a limited sample of the Cyc ontology about domestic objects.

As mentioned in [41], there is no benchmark test available for this kind of task on both human participants and artificial systems. Therefore, we tested our system by comparing our results with the ones described by [41] (table 5, p.23) for the OROC system (to the best of our knowledge, the only available in the literature) by considering the 5 goals they used as testbed. In particular, we asked our system to combine objects in order to obtain the following goals:

$$\mathcal{G}_1 = \{Object, Cutting, Graspable\},$$

$$\mathcal{G}_2 = \{Object, Graspable, LaunchingObjectsAtDistance\},$$

$$\mathcal{G}_3 = \{Object, Support, LiftingFromTheGround\},$$

$$\mathcal{G}_4 = \{CandlewithSupport\},$$

$$\mathcal{G}_5 = \{Notebook\}$$

In particular, we discarded the goals 4 and 5 since they are intended as a composition based on a simple meronymy. Goal 4, in fact, is achievable by just composing the two objects *Candle* and *Candle Support* available in the knowledge base. Also the goal of realizing a Notebook was achievable by composing two constituents part of the object available in the KB: *Blank pages* and *Cover*. Such goals can be easily reached by using a standard Description Logic reasoner, without resorting to the sophistication of $\mathbf{T}^{\text{cl}}$ for the commonsense conceptual composition. For the goals 1, 2 and 3, on the contrary we adopted the framework proposed in this paper.

As mentioned, we have considered an extension of the knowledge base `Open Cyc` where we manually introduced, in the language of $\mathbf{T}^{\text{cl}}$, typicality-based properties/inclusions that were not originally available in the ontology due to the fact that standard ontological semantics does not support representing and reasoning on typicality and exceptions [14, 11]. An example of the introduced inclusions/properties (for the concepts Shelf, Stone, Stump, RubberBand) is reported below:

Shelf $\sqsubseteq$ *Object*

0.8 :: $\mathbf{T}(\textit{Shelf}) \sqsubseteq \textit{Wood}$

0.9 :: $\mathbf{T}(\textit{Shelf}) \sqsubseteq \textit{Rectangular}$

0.8 :: $\mathbf{T}(\textit{Shelf}) \sqsubseteq \textit{Containment}$

0.8 :: $\mathbf{T}(\textit{Shelf}) \sqsubseteq \textit{Support}$

0.8 :: $\mathbf{T}(\textit{Stump}) \sqsubseteq \textit{Wood}$

0.7 :: $\mathbf{T}(\textit{Stump}) \sqsubseteq \textit{Medium}$

0.8 :: $\mathbf{T}(\textit{Stump}) \sqsubseteq \textit{Linear}$

0.7 :: $\mathbf{T}(\textit{Stump}) \sqsubseteq \textit{LiftingFromGround}$

0.7 :: $\mathbf{T}(\textit{Stump}) \sqsubseteq \textit{Support}$

Stone $\sqsubseteq$ *MineralAggregate*

0.7 :: $\mathbf{T}(\textit{Stone}) \sqsubseteq \textit{Roundish}$

0.7 :: $\mathbf{T}(\textit{Stone}) \sqsubseteq \textit{Greyish}$

0.7 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{BuildingArrowHeads}$

0.8 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{ShapingObjects}$

0.7 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{Cutting}$

0.6 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{Support}$

0.8 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{StrikeAtDistance}$

0.9 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{Graspable}$

0.7 :: $\mathbf{T}(\text{Stone}) \sqsubseteq \text{Narrow}$

$\text{RubberBand} \sqsubseteq \text{Object}$

$\text{RubberBand} \sqsubseteq \text{Plastic}$

0.9 :: $\mathbf{T}(\text{RubberBand}) \sqsubseteq \text{Propeller}$

0.9 :: $\mathbf{T}(\text{RubberBand}) \sqsubseteq \text{LaunchingObjectsAtDistance}$

0.7 :: $\mathbf{T}(\text{RubberBand}) \sqsubseteq \text{Small}$

Given a KB extended in $\mathbf{T}^{\text{cl}}$ as reported above, we employed our system for solving the first 3 goals. For what concerns the first goal, i.e. where the purpose of our intelligent system consisted in looking for a graspable object able to cut, the system was not able to find a unique object satisfying all the properties and, therefore, proposed the combination $\text{Stone} \sqcap \text{Branch}$ as a solution, thus suggesting a combined concept having the characteristics resembling a rudimentary *KnifeWithAWoodHandle*.

For what concerns the second goal, where the system was asked to look for a graspable object able to launch objects at distance, the system combined the concepts *Branch* and *RubberBand*, being those with the highest rank with respect to $\mathcal{G}_2$. The $(\text{Branch} \sqcap \text{RubberBand})$ -revised knowledge base, suggested by adopting *Branch* as the HEAD, is such that all the properties of both concepts are considered, with the exception of *Support*. Therefore the knowledge base of the agent is extended (among the others) by the following inclusions:

0.9 :: $\mathbf{T}(\text{Branch} \sqcap \text{RubberBand}) \sqsubseteq \text{Graspable}$

0.9 :: $\mathbf{T}(\text{Branch} \sqcap \text{RubberBand}) \sqsubseteq \text{LaunchingObjectsAtDistance}$

and the combination *Branch* $\sqcap$ *RubberBand* is a solution for the goal $\mathcal{G}_2$. The intentional description of the combined concept for $\mathcal{G}_2$ corresponds to the concept *Slingshot*.

For what concerns the third goal, the system provides a solution by combining *Shelf* and *Stump*. The intentional description of the combined concept for $\mathcal{G}_3$ corresponds to the concept *Table*.

Therefore our system provided the same results provided in the OROC system. In addition with respect to the work by [41] (that, on the other hand, collected human data for the task of object substitution and not for the one of object composition), we collected data from 36 human subjects that were asked to solve the same type of goal by considering the same subset of domestic object considered by our system for the combination. The results provided for the 3 goals are reported in Figure 1. In particular, the most rated results are compliant with the results reported by both our system and by OROC. Apart from the mere choice of the concept to select for the combination, we also asked to the human subject to indicate which kind of object they were thinking for justifying their combination (the datum is reported in the round parenthesis in the table, along with the percentage of the people that responded in favor of the most rated combination). Interestingly enough, human subjects were also able to provide multiple valid solutions for these constrained goals. In particular, for the $\mathcal{G}_1$ they also provided as solution the combination of the concepts *Stone* and *Towel* (justified as the possibility of having a sort of soft material through which to handle a stone). For $\mathcal{G}_2$ the second most rated answer was *RubberBand* and *Towel* (justified with the idea of passing the towel in the circular rubberband and using it as a sort of soft support for the rubberband) while for $\mathcal{G}_3$ the second solution consisted in the combination of the concepts *Vase* with *Shelf* (where the supporting element was constituted by the vase). To test whether the system was able to provide similar results, we repeated the task by deleting the originally obtained combination and by looking at its the second ranked proposal (if any). For $\mathcal{G}_1$ the system was not able to find any suitable combination, for $\mathcal{G}_2$ it proposed the combination of the concepts *RubberBand* and *Book* (with the latter object used a rigid support) and for $\mathcal{G}_3$ the system proposed to combine the concepts *Stump* and *SurfingBoard* (by using the latter element as a support). The figure 1 also

reports the second most rated results from both human subjects and our system ⁸).

	$\mathcal{G}_1$	$\mathcal{G}_2$	$\mathcal{G}_3$
System	<i>Stone</i> $\sqcap$ <i>Branch</i>	<i>Branch</i> $\sqcap$ <i>RubberBand</i>	<i>Shelf</i> $\sqcap$ <i>Stump</i>
Human	<i>Stone</i> $\sqcap$ <i>Branch</i> (<i>KnifeWithHandle</i> , 52%)	<i>Branch</i> $\sqcap$ <i>RubberBand</i> (<i>Slingshot</i> , 42%)	<i>Shelf</i> $\sqcap$ <i>Stump</i> (<i>Table</i> , 59%)
System	-	<i>Book</i> $\sqcap$ <i>RubberBand</i>	<i>Stump</i> $\sqcap$ <i>SurfBoard</i>
Human	<i>Stone</i> $\sqcap$ <i>Towel</i> (13, 3%)	<i>Towel</i> $\sqcap$ <i>RubberBand</i> (10, 8%)	<i>Vase</i> $\sqcap$ <i>Shelf</i> (22, 5%)

Figure 1: Comparison on Concept Composition in a Domestic Domain.

6. Beyond Subgoalting: Compliance and Extension of SOAR procedures

In this section we show how the proposed system can integrate, and extend, the classical subgoalting mechanism embedded in a cognitive architecture like SOAR [20] (with a particular reference to the *impasse* mechanisms developed in such architecture). In our opinion, this compliance represents an important aspect to point out since SOAR is one of the most mature cognitive architectures and has been used by many researchers worldwide during the last 30 years in the field of cognitive modelling and intelligent systems. This system was considered by Allen Newell a candidate for a Unified Theory of Cognition [39] and still represents an important pillar in the effort of building a general integrated model of cognition [21]. This system adheres strictly to Newell and Simon’s physical symbol system hypothesis [40] which states that symbolic processing is a necessary and sufficient condition for intelligent behavior. One of the main themes in SOAR is that all cognitive tasks can be represented by problem spaces that are searched by production rules grouped into operators. These production rules are fired in parallel to produce reasoning cycles. From a representational perspective, SOAR exploits symbolic representations of knowledge (called chunks) and use pattern matching to select relevant knowledge elements. Basically, when a production

⁸Additional results going beyond the 2nd most rated results are not reported since they only regard a minimal percentage of answers.

Figure 2: The SOAR Cognitive Architecture and, in red, the compliance with the proposed system in order to extend the classical subgoaling procedure of the architecture.

rule matches the contents of declarative (working) memory, then the rule fires and the content from the declarative memory (called Semantic Memory in SOAR) is retrieved.

Such type of knowledge structures, however, are usually heavily used to perform standard logical reasoning and, as a consequence, are strongly biased towards a “classical” conceptualisation of knowledge in terms of necessary or sufficient conditions and are not equipped with commonsense representational and reasoning knowledge

components⁹. If a problem (an *impasse* in SOAR terms) arises due to the fact that certain knowledge is lacking, resolving this impasse automatically becomes the new goal (and this process is known as *subgoaling*). This new goal becomes a subgoal of the original one, which means that once the subgoal is achieved, control is returned to the main goal. The subgoal has its own problem space, state and possible set of operators. Whenever the subgoal has been achieved it passes its results to the main goal, thereby resolving the impasse. Learning is keyed to the subgoaling process: whenever a subgoal has been achieved, new knowledge is added to the knowledge base to prevent the impasse that produced the subgoal from occurring again (this learning process is known as *chunking*). If an impasse occurs because the consequences of an operator are unknown, and in the subgoal these consequences are subsequently found, knowledge is added to SOAR’s memory about the consequences of that operator. An important feature in SOAR concerns the fact that it can also use external input as part of its impasse resolution process, therefore new knowledge can extend the Semantic Memory of SOAR and can be incorporated into the learned rules. In this context, the proposed system can be integrated, and can extend, the SOAR subgoaling procedure as illustrated in the figure 2. The process of bi-directional translation between a chunk-like representation and the language of $\mathbf{T}^{\text{cl}}$ can be provided as introduced in [13] and implemented in [17], where a typicality-based KB properties is translated into a standard Description Logic knowledge base (corresponding to a chunk-based symbolic representation in SOAR). In particular, the overall approach is compliant with the idea of a goal-directed contextual activation of concepts obtained via a process of knowledge “proxyfication” [26] from the long-term memory to the short term memory of a cognitive agent (already employed in knowledge-based systems like DUAL-PECCS [33], integrated with different cognitive architectures, including SOAR [34, 32]). A final element emerging from the described compliance consists in the fact that the output of the new subgoaling procedure (i.e. the novel concept dynamically generated in the KB and made available

⁹This problem arises despite the fact that the chunks in SOAR can be represented as a sort of frame-like structures containing some commonsense (e.g. prototypical) information. We remind to [27] for details analysis on such issue.

in the working memory of the architecture to solve the original goal) can be used in the SOAR learning mechanism known as *chunking*, which converts the obtained concept used to solve the goal at hand in the procedural memory of the system in order to avoid to perform *ex-novo* the same reasoning cycle in case the agent encounters again the same goal to solve.

From a more implementative point of view, the above mentioned integration between the Semantic Memory in SOAR (SMEM) and our hybrid KB can be obtained as follows: SMEM is accessed through two dedicated working memory channels, called $\hat{command}$ and $\hat{result}$. In particular, $\hat{command}$ is the branch of the working memory buffer where the GOAL setting takes places. In case the goal cannot be satisfied, this kind of request, instead of launching a standard search in the SOAR SMEM, can use our system to select which concepts can be potentially combined to extend the available knowledge and to solve the goal in hand. This kind of connection can be done by modifying the SOAR kernel and by creating novel RHS (Right Hand Side) functions able to launch our system, and its T^{CL} knowledge base in order to take advantage of the reasoning procedure presented in the above sections. The result of this process will produce in output a novel prototype-based representation that is can be used to solve and goal. Such result can be stored in the $\hat{result}$ channel, the branch of the SOAR working memory buffer devoted to acquiring the output from the external modules. Once the result of our system is “proxyfied” and the goal is solved, it can then be used in the chunking mechanism of the architecture.

7. Discussion and Conclusions

The capability of integrating and generating novel knowledge to solve problems is one of the 14 functional criteria individuated by Allen Newell, in his book *Unified Theories of Cognition* [39], for a cognitive architecture (criterium nr.6). In this paper, we have presented a system aimed at specifically addressing this problem by proposing an extension of classical subgoalting procedures through a dynamical, goal-driven, enrichment of an agent knowledge base obtained via a procedure exploiting a process of commonsense conceptual combination based on the logic T^{CL} .

The proposed approach has been tested in the task of object composition and compared with the available results of the system OROC [41] that is, to the best of our knowledge, the first system proposing a proof-of-concept procedure for the evaluation of such tasks. In particular, we have shown how our framework is able to generate the same results provide by the OROC system by adopting different representational and reasoning assumptions. In addition, we have also compared the obtained results with a preliminary evaluation involving human subjects in the task of object composition. As a further element, we have also shown that the proposed framework is compliant with all the major mechanisms available in the SOAR cognitive architecture and, as such, it can be effectively used to extend its subgoaling procedures (and therefore the reasoning capabilities of the agents equipped with such architecture). In the next section we review the related works and conclude with some pointers to future developments.

7.1. Related Works

For what concerns 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 [25] present a detailed analysis of the limits of the set-theoretic approaches [37], the fuzzy logics [47, 8] (whose limitations was already shown in [42, 45, 19]), the vector-space models [36] and quantum probability approaches [1] proposed to model this phenomenon. In addition, they propose to use hierarchical conceptual spaces [12] 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 [25, 19]: 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 explicitly assumes that the combination is not commutative (i.e. the dif-

ferent attribution of the HEAD-MODIFIER roles does not provide the same combined concept) and that iv) there are dominance effects in the concepts to be combined (both these effects are obtained via the HEAD-MODIFIER heuristics).

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 [38]). In this setting, [7] 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 formalism, which does not allow to reason about typicality and defeasible inheritance. More recently, a different approach is proposed in [9], 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 [31], the approach adopted in our system is flexible enough to be applied also to the case of conceptual blending. There is no evidence, however, that both the frameworks of [7] and [9] would be able to model (in toto or in part) conceptual combination problems like the object composition task. As such, $\mathbf{T}^{\text{cl}}$ 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 $\mathbf{T}^{\text{cl}}$, relying on the prover RAT-OWL [17] 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 Python and exploits a translation of an $\mathcal{ALC} + \mathbf{T}_R$ knowledge base into standard $\mathcal{ALC}$.

7.2. Future Works

In future research we aim at extending our approach to more expressive symbolic formalisms and Description Logics such as, for example, those underlying the standard OWL language. Starting from the work of [15], applying the logic with the typicality

operator and the rational closure to the logic $\mathcal{SHIQ}$, we intend to study whether and how $\mathbf{T}^{\text{CL}}$ could provide an alternative solution to the problem of the “all or nothing” behavior of rational closure with respect to property inheritance. This will allow one to express prototypical properties with a richer language, as well as to perform useful inferences.

The logic $\mathbf{T}^{\text{CL}}$ underlying the system is also able to combine more than two concepts at a time, as well as to involve compound concepts (and not only atomic ones) in a concept combination. We aim at extending our approach in order to also exploit this feature. Moreover, in future works, we plan to consider the case in which the system is able to provide a partial solution, satisfying a proper subset of the initial goals.

The system described in section 4 relies on COCOS, a tool for combining concepts in the logic $\mathbf{T}^{\text{CL}}$. In future research, we aim at studying the application of optimization techniques in [3, 5] in order to improve the efficiency of COCOS and, a consequence, of the proposed goal-driven knowledge generation system.

Finally, we aim at extending the evaluation provided in this paper in two directions: the first one concerns the release of a richer dataset to use for the task of Object Composition for testing both human and artificial creativity (and this will require a truly interdisciplinary effort). The second one goes in the direction of testing our dynamic knowledge generation system on larger knowledge bases. This aspect would require to analyze in more detail heuristic aspects concerning the efficiency about the concept selection and combination.

8. References

- [1] 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.
- [2] Aha, D. W. (2018). Goal reasoning: Foundations, emerging applications, and prospects. *AI Magazine*, 39(2).
- [3] Alberti, M., Bellodi, E., Cota, G., Riguzzi, F., and Zese, R. (2017). cplint on

- SWISH: probabilistic logical inference with a web browser. *Intelligenza Artificiale*, 11(1):47–64.
- [4] Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., and Nardi, D. (2003). *The description logic handbook: Theory, implementation and applications*. Cambridge university press.
- [5] Bellodi, E., Lamma, E., Riguzzi, F., Zese, R., and Cota, G. (2017). A web system for reasoning with probabilistic OWL. *Journal of Software: Practice and Experience*, 47(1):125–142.
- [6] Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103(1-2):347–356.
- [7] 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.
- [8] Dubois, D. and Prade, H. (1997). The three semantics of fuzzy sets. *Fuzzy sets and systems*, 90(2):141–150.
- [9] 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.
- [10] Frixione, M. and Lieto, A. (2012). Representing concepts in formal ontologies: Compositionality vs. typicality effects. *Logic and Logical Philosophy*, 21(4):391–414.
- [11] Frixione, M. and Lieto, A. (2014). Towards an Extended Model of Conceptual Representations in Formal Ontologies: A Typicality-Based Proposal. *Journal of Universal Computer Science*, 20(3):257–276.

- [12] Gärdenfors, P. (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press.
- [13] Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2009). ALC+T: a preferential extension of description logics. *Fundamenta Informaticae*, 96:341–372.
- [14] Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2013). A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, 195:165–202.
- [15] Giordano, L., Gliozzi, V., Olivetti, N., and Pozzato, G. L. (2014). Rational closure in *SHIQ*. In *DL 2014, 27th International Workshop on Description Logics*, volume 1193 of *CEUR Workshop Proceedings*, pages 543–555. CEUR-WS.org.
- [16] 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.
- [17] 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.
- [18] Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15(1):55–71.
- [19] Hampton, J. A. (2011). Conceptual combinations and fuzzy logic. *Concepts and fuzzy logic*, 209.
- [20] Laird, J. (2012). *The Soar cognitive architecture*. MIT Press.
- [21] Laird, J. E., Lebiere, C., and Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *Ai Magazine*, 38(4).

- [22] Lawry, J. and Tang, Y. (2009). Uncertainty modelling for vague concepts: A prototype theory approach. *Artificial Intelligence*, 173(18):1539–1558.
- [23] Lehmann, D. and Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60.
- [24] Lenat, D. B. (1995). Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38.
- [25] Lewis, M. and Lawry, J. (2016). Hierarchical conceptual spaces for concept combination. *Artificial Intelligence*, 237:204–227.
- [26] Lieto, A. (2014). A computational framework for concept representation in cognitive systems and architectures: Concepts as heterogeneous proxytypes. *Procedia Computer Science*, 41:6–14.
- [27] Lieto, A., Lebiere, C., and Oltramari, A. (2018a). The knowledge level in cognitive architectures: Current limitations and possible developments. *Cognitive Systems Research*, 48:39–55.
- [28] Lieto, A., Pozzato, G., and Valesse, A. (2018b). COCOS: a typicality based COnccept COmbination System . In Montali, M. and Felli, P., editors, *Proceedings of the 33rd Italian Conference on Computational Logic (CILC 2018)*, CEUR Workshop Proceedings, pages 55–59, Bozen, Italy.
- [29] Lieto, A. and Pozzato, G. L. (2018). A description logic of typicality for conceptual combination. In Ceci, M., Japkowicz, N., Liu, J., Papadopoulos, G. A., and Ras, Z. W., editors, *Foundations of Intelligent Systems - 24th International Symposium, ISMIS 2018, Limassol, Cyprus, October 29-31, 2018, Proceedings*, volume 11177 of *Lecture Notes in Computer Science*, pages 189–199. Springer.
- [30] Lieto, A. and Pozzato, G. L. (2019a). Applying a description logic of typicality as a generative tool for concept combination in computational creativity. *Intelligenza Artificiale*, (to appear).

- [31] Lieto, A. and Pozzato, G. L. (2019b). A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics. *arXiv preprint arXiv:1811.02366*.
- [32] Lieto, A., Radicioni, D., Rho, V., and Mensa, E. (2017a). Towards a unifying framework for conceptual representation and reasoning in cognitive systems. *Intelligenza Artificiale*, 11(2):139–153.
- [33] 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.
- [34] Lieto, A., Radicioni, D. P., and Rho, V. (2017b). Dual peccs: a cognitive system for conceptual representation and categorization. *Journal of Experimental & Theoretical Artificial Intelligence*, 29(2):433–452.
- [35] Miller, G. A. (1995). Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- [36] Mitchell, J. and Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive science*, 34(8):1388–1429.
- [37] Montague, R. (1973). The proper treatment of quantification in ordinary english. In *Approaches to natural language*, pages 221–242. Springer.
- [38] Nagai, Y. and Taura, T. (2006). Formal description of concept-synthesizing process for creative design. *Design computing and cognition*, pages 443–460.
- [39] Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.
- [40] Newell, A. and Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126.
- [41] Oltețeanu, A.-M. and Falomir, Z. (2016). Object replacement and object composition in a creative cognitive system. towards a computational solver of the alternative uses test. *Cognitive Systems Research*, 39:15–32.

- [42] Osherson, D. N. and Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 9(1):35–58.
- [43] Riguzzi, F., Bellodi, E., Lamma, E., and Zese, R. (2015). 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.
- [44] Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of experimental psychology: General*, 104(3):192.
- [45] Smith, E. E. and Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive science*, 8(4):337–361.
- [46] von Bayern, A. M. P., Danel, S., Auersperg, A., Mioduszevska, B., and Kacelnik, A. (2018). Compound tool construction by new caledonian crows. *Scientific reports*, 8(1):15676.
- [47] Zadeh, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning—i. *Information sciences*, 8(3):199–249.