An Ontology-based Tool for Dynamic Generation, Classification and Recommendation of Novel Contents in Online Libraries

Chiara Barbera[1,†], Antonio Lieto[1,†] and Gian Luca Pozzato[1,*,†]

[1]*Dipartimento di Informatica, Università di Torino, c.so Svizzera 185, 10149 Turin, Italy*

Abstract

In this work we present AMARETTO (dynAMic generAtoR of novEl conTenT in bOoks), an intelligent recommender system exploiting a nonmonotonic extension of Description Logics with typical properties and probabilities to dynamically generate novel contents in Goodreads, the largest website for readers and book recommendations (https://www.goodreads.com). The tool AMARETTO can be used to both the generation/suggestion of novel genres of books and the reclassification of the available items within such new genres. AMARETTO first extracts a prototypical description of the available genres by means of a standard information extraction pipeline, then it generates novel classes of genres as the result of an ontology-based combination of such extracted representations, by exploiting the reasoning capabilities of a probabilistic extension of a Description Logic of typicality. We have tested AMARETTO by reclassifying the available books in Goodreads with respect to the new generated genres, as well as with an evaluation, in the form of a controlled user study experiment, of the feasibility of using the obtained reclassifications as recommended contents. The obtained results are encouraging and pave the way to many possible further improvements and research directions.

Keywords

Cognitive Systems, Recommender System, Knowledge Invention, Description Logics, nonmonotonic reasoning, probabilities, reasoning about typicality

1. Introduction

Dynamic generation of novel knowledge via conceptual recombination is a relevant phenomenon. It highlights some crucial aspects of the knowledge processing capabilities in human cognition. Indeed, such ability concerns high-level capacities associated to creative thinking and problem solving. The recent literature suggests the relevance of this topic [1, 2, 3], however, it still represents an open challenge in the field of artificial intelligence [4]. Indeed, dealing with this problem requires, from an AI perspective, the harmonization of two conflicting requirements: the need of a syntactic and semantic compositionality – typical of logical systems

1st Italian Workshop on Artificial Intelligence for Cultural Heritage (AI4CH22), co-located with the 21st International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022). 28 November 2022, Udine, Italy.

*Corresponding author.

†These authors contributed equally.

chiara.barbera@edu.unito.it (C. Barbera); antonio.lieto@unito.it (A. Lieto); gianluca.pozzato@unito.it (G. L. Pozzato)

https://www.antoniolieto.net (A. Lieto); http://www.di.unito.it/~pozzato/ (G. L. Pozzato)

0000-0002-8323-8764 (A. Lieto); 0000-0002-3952-4624 (G. L. Pozzato)

– and the one concerning the exhibition of typicality effects. Such requirements, however, can be hardly accommodated in standard symbolic systems, including formal ontologies [5]. According to a well-known argument [6], in fact, prototypes, namely 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, a typical fish is grayish, but a typical pet fish is neither furry nor grayish, on the contrary, typically, it is red. The *pet fish* phenomenon is a paradigmatic example of the difficulty to address when building formalisms and systems trying to imitate this combinatorial human ability. Examples of such difficulties concern: handling exceptions to attribute inheritance, handling the possible inconsistencies arising between conflicting properties of the concepts to be combined etc.

In this work, we exploit an ontology-based framework able to account for this type of human-like concept combination and we show how it can be used as a tool for the generation and the suggestion of novel editorial content, following the same idea of DENOTER [7], a Knowledge-Based System for the dynamic generation and classification of novel contents in multimedia broadcasting. In particular, we adopt the recently introduced nonmonotonic extension of Description Logics (from now on DL, see [8]) able to reason about typicality with probabilities and called $\mathbf{T}^{\text{CL}}$ (typicality-based compositional logic) introduced in [9]. Description Logics are a class of decidable fragments of first order logics that are at the base of Ontology Web Language (OWL and OWL 2), used for the realization of computational ontologies. Nowadays, DLs are the most important and widespread symbolic knowledge-representation formalisms [8]. In the logic $\mathbf{T}^{\text{CL}}$, "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 $\mathbf{T}^{\text{CL}}$ one can consistently express exceptions and reason about defeasible inheritance as well.

Typicality inclusions are also equipped by a real number $p \in (0.5, 1]$ representing the probability/degree of belief in such a typical property: this allows us to define a semantics inspired to the DISPONTE semantics [10] characterizing probabilistic extensions of DLs, which in turn is used in order to describe different *scenarios* where only some typicality properties are considered. Given a KB containing the prototypical description of two concepts C_H and C_M occurring in it, we then consider only *some* scenarios in order to define a revised knowledge base, enriched by typical properties of the combined concept $C \sqsubseteq C_H \sqcap C_M$ by also implementing some heuristics coming from the cognitive semantics.

In this work we exploit the logic $\mathbf{T}^{\text{CL}}$ in order to dynamically generate novel knowledge about literary genres and their properties by means of a mechanism for commonsense combination. This generative and creative capacity has been tested in the context of the online platform Goodreads (https://www.Goodreads.com), one of the largest site for readers and book recommendations providing information about literary genres as well as about hundreds of books. We introduce the system AMARETTO (dynAMic generAtoR of novEl conTenT in bOoks) which, first, automatically builds prototypes of existing *basic* genres in Goodreads (Children's, Fantasy, Horror, and so on). In Goodreads, each book is explicitly marked as belonging to one or more basic genres. By means of a web crawler, AMARETTO extracts information about *concepts* or *properties* occurring with the highest frequencies in the textual descriptions of the book available in the online platform. Such prototypes are formalized by means of a $\mathbf{T}^{\text{CL}}$ knowledge base, whose TBox contains both *rigid* inclusions of the form

$$BasicGenre \sqsubseteq Concept,$$

in order to express essential desiderata but also constraints, for instance $Childrens \sqsubseteq Good$ (contents for children must have goods) and $Childrens \sqsubseteq \neg Sex$ (due to law restrictions, sexual contents for kids must be forbidden), as well as *prototypical* properties of the form

$$p \ :: \ \mathbf{T}(BasicGenre) \sqsubseteq TypicalConcept,$$

representing typical concepts of a given genre, where p is a real number in the range $(0.5, 1]$, expressing the frequency of such a concept in items belonging to that genre: for instance, $0.825 \ :: \ \mathbf{T}(Horror) \sqsubseteq House$ is used to express that the typical horror book contains/refers to the concept House with a frequency/probability/degree of belief of the 82.5%, and such a degree is automatically extracted by AMARETTO from the description of the books currently available on Goodreads and classified as belonging to such a genre.

Given the knowledge base with the prototypical descriptions of basic genres, AMARETTO exploits the reasoning capabilities of the logic $\mathbf{T}^{\text{CL}}$ in order to generate new *derived* genres as the result of the creative combination of two basic ones. It is possible to use AMARETTO also for combining derived genres with either basic or derived ones. With the prototypical descriptions of the derived genres at hand, the tool AMARETTO reclassifies the books of Goodreads taking such new, derived genres into account. The basic idea is as follows: a book can be considered as belonging to/is recommended for a given new generated genre if its metadata (name, description, title) contain all the rigid properties as well as at least the 30% of the typical properties of the prototype of such a derived genre.

We have also tested AMARETTO by performing two different kinds of evaluation: on the one hand, an automatic evaluation, on the other hand an evaluation of the satisfaction of users. In both cases, the results that we have obtained seem promising, witnessing that the tool AMARETTO could represent a first step in the direction of designing a novel, data-driven, logic-based, "white box" smart recommender system.

2. Logical Reasoning for Concept Combination: the Description Logic $\mathbf{T}^{\text{CL}}$

The tool AMARETTO exploits the Description Logic $\mathbf{T}^{\text{CL}}$ [9] for the generation of new literary genres as the combination of two existing ones.

The language of $\mathbf{T}^{\text{CL}}$ extends the basic DL $\mathcal{ALC}$ by *typicality inclusions* of the form

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

where $p \in (0.5, 1]$ is a real number representing its degree of belief, whose meaning is that "we believe with degree/probability p that, normally, Cs are also Ds". We avoid probabilities $p \leq 0.5$ since it would be misleading for typicality inclusions, since typical knowledge is known to come with a low degree of uncertainty.

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, and $\mathcal{A}$ is the ABox, i.e. a finite set of formulas of the form either $C(a)$ or $R(a, b)$, where $a, b \in \mathtt{O}$ and $R \in \mathtt{R}$.

The Description Logic $\mathbf{T}^{\mathrm{CL}}$ relies on the DL of typicality $\mathcal{ALC} + \mathbf{T_R}$ introduced in [11], which allows to describe the *protoype* of a concept, in this case a literary genre. 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. The semantics of the $\mathbf{T}$ operator is characterized by the properties of *rational logic* [12], recognized as the core properties of nonmonotonic reasoning. The Description Logic $\mathcal{ALC} + \mathbf{T_R}$ is characterized by a minimal model semantics corresponding to an extension to DLs of a notion of *rational closure* as defined in [12] 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, the operator $\mathbf{T}$ inherits well-established properties like *specificity* and *irrelevance*; in the example, the Description Logic $\mathcal{ALC} + \mathbf{T_R}$ allows one to infer that $\mathbf{T}(\mathit{Student} \sqcap \mathit{Tall}) \sqsubseteq \mathit{Young}$ (being tall is irrelevant with respect to being young) and, if one knows that Annekee is a typical senior student, to infer that she is not young, giving preference to the most specific information.

The Description Logic of typicality with rational closure is finally extended in order to deal with concept combination: the logic $\mathbf{T}^{\mathrm{CL}}$ considers a distributed semantics similar to DISPONTE [13] for probabilistic DLs. This logic allows one to label inclusions $\mathbf{T}(C) \sqsubseteq D$ with a real number between 0.5 and 1, representing its degree of belief, assuming that each axiom is independent from each others. Degrees in typicality inclusions allow to define a probability distribution over *scenarios*: intuitively, a scenario is obtained by choosing, for each typicality inclusion, whether it is considered as true or false. In an extension of the above example, we could have the following KB:

(1) $\mathit{SeniorStudent} \sqsubseteq \mathit{Student}$
(2) $0.7 \ :: \ \mathbf{T}(\mathit{Student}) \sqsubseteq \mathit{Young}$
(3) $0.95 \ :: \ \mathbf{T}(\mathit{SeniorStudent}) \sqsubseteq \neg \mathit{Young}$
(4) $0.85 \ :: \ \mathbf{T}(\mathit{SeniorStudent}) \sqsubseteq \mathit{Married}$

We consider eight different scenarios, representing all possible combinations of typicality inclusion, for instance $\{((2), 1), ((3), 1), ((4), 0)\}$ represents the scenario in which (2) and (3) hold, whereas (4) is not considered. The standard inclusion (1) holds in every scenario, representing a rigid property not admitting exceptions. We equip each scenario with a probability depending on those of the involved inclusions: the scenario of the example has probability $0.7 \times 0.95 \times (1 - 0.85)$, since 2 and 3 are involved, whereas 4 is not. Such probabilities are then taken into account in order to select the most adequate scenario describing the prototype of the combined concept.

In order to capture the ability of combining concepts, the logic $\mathbf{T}^{\mathrm{CL}}$ exploits a method inspired by cognitive semantics [14] 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 $C \sqsubseteq C_H \sqcap C_M$.

Formally, given a KB $\mathscr{K} = \langle \mathscr{R}, \mathscr{T}, \mathscr{A} \rangle$ and given two concepts C_H and C_M occurring in $\mathscr{K}$, the logic $\mathbf{T}^{\text{CL}}$ allows defining a prototype of the concept obtained by the combination of the HEAD C_H and the MODIFIER C_M. Such a prototype contains typical properties of the form

$$p \ :: \ \mathbf{T}(C_H \sqcap C_M) \sqsubseteq D$$

and are obtained by considering blocks of scenarios with the same probability, in decreasing order from the highest to the lowest. We first discard all the inconsistent scenarios, then:

- we discard *trivial* scenarios, that is to say scenario 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 all scenarios inheriting all the properties of the HEAD, including the one where all the properties of both concepts HEAD and MODIFIER, are inherited, are discarded since they 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; as an example, in the combination of the literary genres *Mystery* – as the HEAD – and *Christian* – MODIFIER — the logic could have to deal with the conflicting properties:

  $0.85 \ :: \ \mathbf{T}(\mathit{Horror}) \sqsubseteq \mathit{Killer}$
  $0.60 \ :: \ \mathbf{T}(\mathit{Christian}) \sqsubseteq \neg \mathit{Killer}$

  In this case, all scenarios where the latter inclusion is considered are discarded, even if the former one is discarded too, since they reject the HEAD/MODIFIER heuristics;
- 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 are those selected by the logic $\mathbf{T}^{\text{CL}}$. The ultimate output of our mechanism is a knowledge base whose set of typicality properties is obtained by adding to those of the initial knowledge, the inclusions considered in the selected scenario(s), namely those describing the prototype of the combined concept $C_H \sqcap C_M$. If more than one scenario is selected by the logic $\mathbf{T}^{\text{CL}}$, then several alternative final knowledge bases are provided.

Formally, given a scenario w satisfying the above properties, we define the properties of $C_H \sqcap C_M$ as the set of inclusions $p \ :: \ \mathbf{T}(C_H \sqcap C_M) \sqsubseteq D$, for all $\mathbf{T}(C_H \sqcap C_M) \sqsubseteq D$ 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. $p \ : \ \mathbf{T}(C_H) \sqsubseteq D \in \mathscr{T}$;
- 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. $p \ : \ \mathbf{T}(C_M) \sqsubseteq D \in \mathscr{T}$.

The knowledge base obtained as the result of combining concepts C_H and C_M i is defined as:

$$\mathscr{K}_{C_H \sqcap C_M} = \langle \mathscr{R}, \mathscr{T} \cup \{p \ : \ \mathbf{T}(C_H \sqcap C_M) \sqsubseteq D\}, \mathscr{A} \rangle,$$

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.

It can be shown that reasoning in the logic $\mathbf{T}^{\text{CL}}$ is ExpTime-complete. Since reasoning in standard $\mathscr{ALC}$ is ExpTime-complete too, this means that, even if we add probabilities, typicalities, and scenarios for combining prototypes, we remain in the same complexity class of the initial, monotonic Description Logic.

3. The tool AMARETTO: Automatic Generation of Novel Literary Genres

We describe the tool AMARETTO, the system exploiting the logic $\mathbf{T}^{\text{CL}}$ in order to generate and suggest novel literary genres for Goodreads (https://www.Goodreads.it). AMARETTO is implemented in Python and it makes use of the library `owlready2` (https://pythonhosted.org/Owlready2/) for relying on the services of efficient DL reasoners (like HermiT).

As already sketched, the tool AMARETTO first builds a prototypical description of basic literary genres available in Goodreads. At present, in order to provide a first evaluation of the system, we consider the following literary genres: Children's, Fantasy, Fiction, History, Horror, Mystery, Romance, and Thriller. To this aim, a web crawler extracts metadata from books and texts available on the platform. More in detail, for each item the crawler extracts (i) the genre to which it belongs and (ii) the set of "significant" words (i.e., excluding prepositions, proper names, articles, etc.) occurring in the description of each item, as well as their frequency. These information are used in order to provide a description of each basic genre in terms of its typical properties in the logic $\mathbf{T}^{\text{CL}}$, where the frequency of a concept/word for a genre is obtained from the number of occurrences of such a concept/word in the items belonging to that genre. The five properties with the highest frequency over 0.5 are included in the prototypical description of each basic genre. Formally, similarly to what done in [7, 15, 16], we have:

Definition 1. Given a book b, let $\mathscr{S}_b$ be the set of significant concepts extracted for b by the web crawler, and let $Concept \in \mathscr{S}_b$. Let $n_{b,Concept}$ be the number of occurrences of *Concept* in the description of b. We define the frequency $f_{b,Concept}$ of concept *Concept* for the item b as

$$f_{b,Concept} = \frac{n_{b,Concept}}{\sum_{D \in \mathscr{S}_b} n_{b,D}}.$$

Definition 2. Given a basic literary genre *Genre*, let $\mathscr{B}$ be the set of books belonging to *Genre*, and let $\mathscr{S}_{Genre}$ be the set of concepts occurring in such items, i.e. $\mathscr{S}_{Genre} = \bigcup_{b \in \mathscr{B}} \mathscr{S}_b$, where $\mathscr{S}_b$.

Given a concept $Concept \in \mathscr{S}_{Genre}$ and a book $b \in \mathscr{B}$, let $n_{b,Concept}$ be the number of occurrences of *Concept* in the description of b. We define $n_{Genre,Concept}$ the number of occurrences of *Concept* in the description of items of *Genre*, i.e.

$$n_{Genre,Concept} = \sum_{b \in \mathscr{B}} n_{b,Concept}.$$

We also define the frequency of *Concept* for a genre *Genre*, written $f_{Genre,Concept}$:

$$f_{Genre,Concept} = \frac{n_{Genre,Concept}}{\sum_{C \in \mathcal{S}_{Genre}} n_{Genre,C}}.$$

The prototypical description of a basic literary *Genre* in$\mathbf{T}^{\text{CL}}$ is defined as the set of inclusions

$$\begin{array}{lll} p_1 & :: & \mathbf{T}(Genre) \sqsubseteq TypicalConcept_1 \\ p_2 & :: & \mathbf{T}(Genre) \sqsubseteq TypicalConcept_2 \\ \vdots & & \\ p_5 & :: & \mathbf{T}(Genre) \sqsubseteq TypicalConcept_5, \end{array}$$

where $TypicalConcept_1$, $TypicalConcept_2$, ..., $TypicalConcept_5$ are the five concepts in $\mathcal{S}_{Genre}$ with the highest frequencies higher than 50%; frequencies are then also used as degrees of belief of the respective inclusions. If needed, in some cases we have also manually added some rigid properties, thus integrating the bottom-up, data-driven, process of prototype formation with top down expert knowledge. Recalling the example in the Introduction, we have exploited this opportunity given by the logic $\mathbf{T}^{\text{CL}}$ for imposing some constraints, for instance to avoid that books for children contain/refer to sex and violence. Therefore, the knowledge base generated by the crawler will contain, among typical properties, also rigid inclusions like $Childrens \sqsubseteq \neg Sex$ and $Childrens \sqsubseteq \neg Violence$. The tool AMARETTO generates novel hybrid literary genres by combining existing ones by exploiting the reasoning mechanism provided by the logic $\mathbf{T}^{\text{CL}}$ and described in the previous section. As an example, consider the following prototypes of basic genres *Childrens* and *Mystery*:

$$\begin{array}{lll} Childrens \sqsubseteq \neg Murder & 0.90 :: \mathbf{T}(Childrens) \sqsubseteq Story & Mystery \sqsubseteq Search \\ Childrens \sqsubseteq \neg Blood & 0.85 :: \mathbf{T}(Childrens) \sqsubseteq Children & 0.90 :: \mathbf{T}(Mystery) \sqsubseteq Murder \\ Childrens \sqsubseteq \neg Sex & 0.80 :: \mathbf{T}(Childrens) \sqsubseteq Book & 0.88 :: \mathbf{T}(Mystery) \sqsubseteq Case \\ Childrens \sqsubseteq \neg Evil & 0.80 :: \mathbf{T}(Childrens) \sqsubseteq Readers & 0.85 :: \mathbf{T}(Mystery) \sqsubseteq Killer \\ Childrens \sqsubseteq \neg Violence & 0.78 :: \mathbf{T}(Childrens) \sqsubseteq Life & 0.81 :: \mathbf{T}(Mystery) \sqsubseteq Years \\ Childrens \sqsubseteq Good & & 0.78 :: \mathbf{T}(Mystery) \sqsubseteq Time \end{array}$$

AMARETTO combines the two basic genres by means of a variant of CoCoS [17], a Python implementation of reasoning services for the logic $\mathbf{T}^{\text{CL}}$ in order to exploit efficient DLs reasoners for checking both the consistency of each generated scenario and the existence of conflicts among properties. As an example, the new, derived literary genre obtained by combining Children's and Mystery, with the number of the inherited properties fixed to be no higher than eight, has the following $\mathbf{T}^{\text{CL}}$ description:

$$\begin{array}{ll} 0.90 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Story & 0.90 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Murder \\ 0.85 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Children & 0.88 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Case \\ 0.80 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Book & 0.85 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Killer \\ 0.90 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Readers & 0.78 :: \mathbf{T}(Childrens \sqcap Mystery) \sqsubseteq Time \end{array}$$

Obviously, rigid properties of both basic concepts *Childrens* and *Mystery* are inherited by the derived concept as well.

4. The tool AMARETTO: Reclassifying and Recommending Books with Respect to Generated Literary Genres

Apart from the process of automatic knowledge generation, AMARETTO is also able to reclassify books of Goodreads within the novel derived genres (generated as described in the previous section). As mentioned, indeed, each book is equipped by some information available in Goodreads, including title and description of the item. AMARETTO extracts such information and then computes the frequencies of concepts in it as in Definition 2, in order to compare them with the properties of a derived genre. If the book contains all the rigid properties and at least the 30% of the typical properties of the genre under consideration, then the book is classified as belonging to it. Last, AMARETTO suggests the set of classified contents, in a descending order of compatibility, where a rank of compatibility of a single book with respect to a genre is intuitively obtained as the sum of the frequencies of "compatible" concepts, i.e. concepts belonging to both the book and the prototypical description of the literary genre. Formally:

Definition 3. Given a book b, let *DerivedGenre* be a derived literary genre as defined in Section 3 and let $\mathcal{S}_b$ be the set of concepts/words occurring in b as in Definition 2. Given a knowledge base KB of literary genres built by AMARETTO, we say that b is compatible with *DerivedGenre* if the following conditions hold:

1. b contains all rigid properties of *DerivedGenre*, i.e. $\{C \mid \mathit{DerivedGenre} \sqsubseteq C \in \text{ KB}\} \subseteq \mathcal{S}_b$
2. b contains at least the 30% of typical properties of *DerivedGenre* , i.e.

$$\frac{\mid \mathcal{S}_b \cap \mathcal{S}_{\mathit{DerivedGenre}} \mid}{\mid \mathcal{S}_{\mathit{DerivedGenre}} \mid} \geq 0.3,$$

   where $\mathcal{S}_{\mathit{DerivedGenre}}$ is the set of typical properties of *DerivedGenre* as in Definition 4.

As an example, consider the above derived genre *Childrens* $\sqcap$ *Mystery*, and the book "The Phantom Tollbooth" (https://www.goodreads.com/book/show/378.The_Phantom_Tollbooth?from_search=true&from_srp=true&qid=x8ixl8TNqL&rank=1). It is reclassified in the novel, generated genre *Childrens* $\sqcap$ *Mystery*. Indeed:

- all rigid properties of both basic genres are satisfied, that is to say the only rigid property *Search* belongs to the properties of the book under consideration, whereas neither *Murder* nor *Blood* nor *Sex* nor *Evil* nor *Violence* belong to the properties extracted by the crawler for the book itself;
- more than the 30% of the typical properties of the derived genre are satisfied by the items, in particular "The Phantom Tollbooth" has *Story* with frequency 0.90, *Readers* (0.90), and *Time* (0.90) with a total score of 3.6.

In conclusion, the tool AMARETTO will recommended the song "The Phantom Tollbooth". As usual, recommended songs are listed in a non increasing order by their score.

5. Tests and Evaluations

We have evaluated AMARETTO in order to check whether it could be considered a promising approach for generating new knowledge about literary genres and recommending books in the context of the Goodreads platform.

The first evaluation is completely automatic and inheres the capability of the system of generating novel hybrid literary genres that are able to be populated by the original content of the Goodreads platform via a re-classification mechanism involving the books of the platform. In this case, the success criterion concerns the avoidance of the creation of empty boxes corresponding to the new generated combined genres. With the only exceptions of three genres, all the derived literary genres contain books that are classified under such a genre. It is worth noticing that, in some cases, a book is classified in the derived genre of two basic ones, without being member of the initial genres. For instance, the book "The Cuckoo's Calling" by J.K.Rowling (under the name Robert Galbraith) is classified as belonging to the genre Mystery, however it is re-classified in the genre derived as the combination of Childrens and Fantasy. Similarly, the above mentioned book "The Phantom Tollbooth" by Norton Juster belongs to the genre Fantasy but it is also classified in the combination of Childrens and Mystery.

A second evaluation aimed at measuring the satisfaction of the potential users of the platform when exposed to the contents of the novel categories suggested by AMARETTO. It consisted in a user study involving 20 persons (10 females, 10 males, aged 20-45) that evaluated recommendations generated by the system for seven combined literary genres. All the participants were selected from the same population, i.e., voluntary persons using an availability sampling strategy. Participants were all naive to the experimental procedure and to the aims of the study. This is one of the most commonly used methodology for the evaluation of recommender systems based on controlled small groups analysis [18]. This evaluation was carried out as a classical "*one to one*" lab controlled experiment, that is to say one person at time with one expert interviewer, and we adopted a thinking aloud protocol: this technique consists in recording the verbal explanations provided by the people while executing a given laboratory task. It has been used in the AI literature since the pioneering work by Newell and Simon, as a source to individuate the heuristics used by humans to solve a given task [19, 20]. At this stage, this solution was methodologically preferred with respect to the adoption of large scale online surveys since it allowed us to have more control on the type of thoughts and considerations emerging during the evaluation of the results. In this setting, the users had to start the interview by indicating a couple of preferred genres among those available in Goodreads.

This selection triggered both the activation of a novel hybrid prototypical literary genre by AMARETTO and the corresponding reclassification of books of Goodreads based on such selection. The output of the system, pruned to show the top 5 best results, was then evaluated with a 1-10 voting scale expressing the satisfaction of the received recommendations. The results of this second evaluation are shown in Table 1 and are promising. For each pair, we have reported only one combination (and not the symmetric one, obtained by exchanging the roles of HEAD and MODIFIER), since we have found very few differences between the two alternatives. We have obtained an average rate of 6.67 in a 1-10 scale: this result seems to be promising, especially considering that only some literary genres have been involved in this test. We are currently working on extending our test including all the genres available on GoodReads.

Derived genre	Average mark
Childrens-Fantasy	6,25
Childrens-Mystery	7,17
Fantasy-Horror	6,5
Fantasy-Romance	6,5
History-Romance	7,17
History-Fiction	6,96
Romance-Mystery	6,17
	6,67

Figure 1: Average marks for all the combinations of literary genres.

6. Conclusions and Future Works

In this work we have presented AMARETTO, a knowledge-based system for the dynamic generation of novel literary genres, exploiting the reasoning mechanism of the logic $\mathbf{T}^{\mathrm{CL}}$ in order to generate, reclassify and suggest books as belonging to novel literary genres in the context of Goodreads, one of the world's largest portal for readers and book recommendations. The system has been tested in twofold evaluation showing promising results for both the automatic evaluation and the user acceptability of the recommended books.

AMARETTO exploits the logic $\mathbf{T}^{\mathrm{CL}}$, introduced in order to tackle the task of modelling 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 to this aim. Other attempts similar to the one adopted 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. In this setting, [3] proposed a mechanism for conceptual blending based on the DL $\mathscr{EL}^{++}$. They construct the generic space of two concepts by introducing an upward refinement operator that is used for finding common generalizations of $\mathscr{EL}^{++}$ 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 [1], 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. There is no evidence, however, that both the frameworks of [3] and [1] would be able to model.

The core component of the system AMARETTO relies on CoCoS. In future research, we aim at studying the application of optimization techniques in [21] in order to improve its efficiency and, a consequence, the one of the proposed knowledge generation system.

Furthermore, we aim at extending the evaluation provided in this paper in two directions. The first one concerns the inclusion of all literary genres available in GoodReads. The second one goes in the direction of testing our tool AMARETTO including a suitable extension for the automated extraction of negated typical properties of the form $p\ ::\ \mathbf{T}(Genre) \sqsubseteq \neg P$. This aspect would require to analyze in more detail heuristic aspects concerning the efficiency about the concept selection and combination.

References

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

[2] M. Lewis, J. Lawry, Hierarchical conceptual spaces for concept combination, Artificial Intelligence 237 (2016) 204–227.

[3] R. Confalonieri, M. Schorlemmer, O. Kutz, R. Peñaloza, E. Plaza, M. Eppe, Conceptual blending in EL++, in: Proc. of the 29th Int. Work. on Description Logics, DL 2016, 2016. URL: http://ceur-ws.org/Vol-1577/paper_8.pdf.

[4] M. A. Boden, Creativity and artificial intelligence, Artificial Intelligence 103 (1998) 347–356.

[5] M. Frixione, A. Lieto, Representing concepts in formal ontologies: Compositionality vs. typicality effects, Logic and Logical Philosophy 21 (2012) 391–414.

[6] D. N. Osherson, E. E. Smith, On the adequacy of prototype theory as a theory of concepts, Cognition 9 (1981) 35–58.

[7] E. Chiodino, D. Di Luccio, A. Lieto, A. Messina, G. L. Pozzato, D. Rubinetti, A knowledge-based system for the dynamic generation and classification of novel contents in multimedia broadcasting, in: G. De Giacomo, A. Catalá, B. Dilkina, M. Milano, S. Barro, A. Bugarín, J. Lang (Eds.), ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020, volume 325 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2020, pp. 680–687. URL: https://doi.org/10.3233/FAIA325. doi:10.3233/FAIA200154.

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

[9] A. Lieto, G. L. Pozzato, A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics, Journal of Experimental and Theoretical Artificial Intelligence 32 (2020) 769–804. URL: https://doi.org/10.1080/0952813X.2019.1672799. doi:10.1080/0952813X.2019.1672799.

[10] F. Riguzzi, E. Bellodi, E. Lamma, R. Zese, Probabilistic description logics under the distribution semantics, Semantic Web 6 (2015) 477–501. URL: http://dx.doi.org/10.3233/SW-140154. doi:10.3233/SW-140154.

[11] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, Semantic characterization of Rational Closure: from Propositional Logic to Description Logics, Artificial Intelligence 226 (2015) 1–33. doi:10.1016/j.artint.2015.05.001.

[12] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial Intelligence 55 (1992) 1–60. doi:http://dx.doi.org/10.1016/0004-3702(92)90041-U.

[13] F. Riguzzi, E. Bellodi, E. Lamma, R. Zese, Reasoning with probabilistic ontologies, in: Q. Yang, M. Wooldridge (Eds.), Proceedings of IJCAI 2015, AAAI Press, 2015, pp. 4310–4316. URL: http://ijcai.org/proceedings/2015.

[14] J. A. Hampton, Inheritance of attributes in natural concept conjunctions, Memory & Cognition 15 (1987) 55–71.

[15] A. Lieto, G. L. Pozzato, M. Striani, S. Zoia, R. Damiano, Degari 2.0: A diversity-seeking, explainable, and affective art recommender for social inclusion, Cognitive Systems Research 77 (2023) 1–17. URL: https://www.sciencedirect.com/science/article/pii/S1389041722000456.

doi:https://doi.org/10.1016/j.cogsys.2022.10.001.
[16] A. Lieto, G. L. Pozzato, S. Zoia, V. Patti, R. Damiano, A commonsense reasoning framework for explanatory emotion attribution, generation and re-classification, Knowledge Based Systems 227 (2021) 107166. URL: https://doi.org/10.1016/j.knosys.2021.107166. doi:10.1016/j.knosys.2021.107166.
[17] A. Lieto, G. L. Pozzato, A. Valese, COCOS: a typicality based concept combination system, in: P. Felli, M. Montali (Eds.), Proceedings of the 33rd Italian Conference on Computational Logic, Bolzano, Italy, September 20-22, 2018, volume 2214 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 55–59. URL: http://ceur-ws.org/Vol-2214/paper6.pdf.
[18] G. Shani, A. Gunawardana, Evaluating recommendation systems, in: Recommender systems handbook, Springer, 2011, pp. 257–297.
[19] A. Newell, J. C. Shaw, H. A. Simon, Report on a general problem solving program, in: IFIP congress, volume 256, Pittsburgh, PA, 1959, p. 64.
[20] A. Newell, H. A. Simon, Human problem solving, volume 104, n. 9, Prentice-Hall Englewood Cliffs, NJ, 1972.
[21] M. Alberti, E. Bellodi, G. Cota, F. Riguzzi, R. Zese, cplint on SWISH: probabilistic logical inference with a web browser, Intelligenza Artificiale 11 (2017) 47–64. URL: https://doi.org/10.3233/IA-170106. doi:10.3233/IA-170106.