Towards Understanding Classification and Identification*

Mattia Fumagalli, Gábor Bella, and Fausto Giunchiglia

DISI - Department of Information Engineering and Computer Science,
University of Trento, Italy
{mattia.fumagalli,gabor.bella,fausto.giunchiglia}@unitn.it

Abstract. The paper focuses on two pivotal cognitive functions of both natural and AI agents, namely *classification* and *identification.* Inspired from the theory of teleosemantics, itself based on neuroscientific results, we show that these two functions are complementary and rely on distinct forms of knowledge representation. We provide a new perspective on well-known AI techniques by categorising them as either classificational or identificational. Our proposed *Teleo-KR architecture* provides a high-level framework for combining the two functions within a single AI system. As validation and demonstration on a concrete application, we provide experiments on the large-scale reuse of classificational (ontological) knowledge for the purposes of learning-based schema identification.

Keywords: classification · identification · teleosemantics · cognitive architecture · knowledge representation

1 Introduction

Class and *classification* are powerful notions in computer science and AI, yet the terms hide a variety of interpretations. Library classifications, for instance, are a traditional form of *knowledge organisation* that apply principled methods to structuring written human knowledge. The notion of class as used in *ontologies* by the Semantic Web community, while also a form of knowledge organisation, is different as it is defined through formal logic and it aims to cater to computational applications such as reasoning or data integration. The *machine learning* community also heavily relies on the notion of classification, understanding it as the sorting of a discrete number of input elements into a discrete number of output categories, *classes* or *clusters*, in a supervised or unsupervised manner.

Our paper looks behind the diverse uses of these notions by various AI communities to find that they are not merely the result of different procedural approaches towards similar goals. Rather, they are complementary and serve tasks with markedly different purposes and representational needs.

* This paper was partly supported by the *InteropEHRate* project, co-funded by the European Union (EU) Horizon 2020 programme under grant number 826106.

The theoretical underpinning of our work is the philosophical theory of *teleosemantics* (also known as *biosemantics* or the *teleological theory of mental content*), and in particular Ruth Millikan's results. Teleosemantics is one of the most popular naturalistic explanations of mental representations: it binds together models of cognition, such as the *classical* and *connectionist* models, and has yielded results in fields such as communication theory or genetics [18]. Based on neuroscientific evidence from animals and humans—and thus formulated in total independence from results in AI or computer science—teleosemantics states that *classification* and *identification* are two distinct tasks that are performed using separate devices of the brain that rely on separate representations of knowledge [23].

The paper offers four main contributions. (1) Based on a teleosemantic perspective, we interpret the notions of classification and identification and clarify the difference between the two. Our goal is not to redefine terminology already in use in various fields of AI, but rather to propose a both theoretically and practically useful distinction between kinds of functions that are often conflated into the same task. (2) We categorise a wide range of AI solutions as based on either of the two or their combination, shedding light on why 'classificational' and 'identificational' tasks need different representations of knowledge in order to be efficient. (3) We introduce a novel *Teleo-KR architecture* that bridges these two fundamental cognitive functions and combines them into a unified AI agent. This high-level theoretical framework may serve, in our view, as a blueprint for future hybrid AI solutions for learning to map between different kinds of representations. (4) Finally, we demonstrate the application of the framework on the AI task of matching data schemas via a combined use of the two kinds of knowledge. We implement the setup as a series of experiments on large sets of data schemas and interpret the results.

In section 2, we define and describe classification and identification based on results from teleosemantics. In section 3, we situate well-known AI tasks with respect to these two functions. Section 4 presents the *Teleo-KR architecture* that models cognitive abilities of artificial agents. Section 5 presents our case study on schema identification. Finally, in section 6 we look at the significance of our results and possible future work.

2 Classification and Identification

Teleosemantics considers biological *perceptual-cognitive systems* (PCS)—i.e., what is able to perceive the external environment, to organize sensory information and to know—to be composed of *devices* having specific *functions*. A device corresponds to a biological component of the brain while the notion of function, as used in neurobiology, describes the role fulfilled by the device. Devices perform tasks with specific goals, in relation to other devices or to the external environment.

In a classic clarifying example [21], bees can be considered as PCSs, i.e., *sender/receiver representational systems*, having a device whose function is to accumulate information about a portion of the environment, such as the location where nectar can be found, as well as a device to *communicate* it to other bees, e.g., through the *bee dance*.

'Communication' and 'accumulation' can be generalised as pivotal applications of *classification* and *identification*, respectively. Classificational representations are views over the stream of diverse data sources collected by the representational system over time. Different individuals may have different classificational representations for the same world state (e.g., a car dealer and a mechanic may classify cars differently), and even the same individual may describe the same world state differently according to context and pragmatic requirements. Nevertheless, classes within individual classifications aim to remain consistent and unequivocal.

Identification, in turn, is required to make learning possible: its purpose is to keep track of things over time, to understand whether they were previously encountered or not, and to focus on new incoming information. In contrast to classification, identification relies on an open and adaptive space for diverse, potentially fuzzy, or contradicting information. Identificational representations afford non-invariant knowledge, adapting to changes in how one perceives things over multiple encounters [3, 23].

The device implementing identification relies on knowledge necessary to recognise what is encountered through sensory experience (directly observing the world through seeing, hearing, etc.) and to gather information about it. The device implementing classification builds unequivocal and shareable knowledge from the stream of diverse data collected over experience. Accordingly, a central statement of teleosemantics-which this paper applies to AI as a key contribution-is that *devices may provide their own distinct representations of the world, rather than sharing one common representation. In particular, classificational and identificational representations of knowledge are distinct and are organised in different ways* [20].

Applying these insights to computational agents, we model *identificational representation* (KR^I) as follows:

$$KR^I = \langle S, C^I, \{(s, c^I)\}\rangle \tag{1}$$

where s is a formalization of a *perceptual state*, i.e., a cognitive representation posterior to perception, also called *neural state* in [Barsalou, 1999]). A perceptual state is the initial cognitive encoding of an object encountered by the agent in the external environment. S is a set of all such perceptual states represented within KR^I. c^I is the representational unit of KR^I that [20] calls substance concept and defines as *'nodes that help in storing knowledge and information arriving at the sensory surfaces'* [23]. For our purposes, c^I is a symbol in KR^I that groups perceptual states together as being from the same object in the external environment [11]. C^I is the set of such substance concepts in KR^I. While the simple formalization above suits the purposes of our paper, in practice we expect KR^I to be more complex and fine-grained both for biological and artificial agents.

We model *classificational representation* (KR^C) as follows:

$$KR^C = \langle S, C^C, \{(i, c^C)\}\rangle \tag{2}$$

where $i \in I$ are *instances*, i.e., *representations of occurrences* of a *given object* [14], and $c^C \in C^C$ are *classes*. Here we commit on the classical definition of class provided in [1], taken as a *set of instances*. A true classificational KR may be a superset of this minimal modelling, e.g., based on first-order logic.

It is important to notice that the difference between instance and class (e.g., *'cat'* and *'my cat Misty'*), pivotal in classificational knowledge, does not occur in the identificational knowledge: both always map into a substance concept [14].

A teleosemantic cognitive device can be modelled as a pair consisting of a knowledge representation and a cognitive function: $D = \langle KR^D, f^D\rangle$. Accordingly, the classification and identification devices are composed, respectively, as $D^C = \langle KR^C, f^C\rangle$ and $D^I = \langle KR^I, f^I\rangle$, where f^C and f^I correspond to the cognitive functions of classification and identification.

We model *identification* (f^I) as the function:

$$f^I : \langle S, KR^I\rangle \to C^I \tag{3}$$

that assigns perceptual states resulting from an encounter to a given substance concept. For example, recognising a black shape on a photo as *'a cat'* or *'my cat Misty'* is an act of identification.

We model *classification* (f^C) as the function:

$$f^C : \langle I, KR^C\rangle \to C^C \tag{4}$$

that assigns the instances of a given classification to a given class. The statement *'cats are mammals'*, where the *mammal* is applied to *cat*, both defined within KR^C, is an example of classification.

The representation of identificational knowledge via the *substance concept* strongly relates, in our view, to what in cognitive linguistics is called *basic level category*. As shown in Eleanor Rosch's experiments, the power of identifying something (such as *a cat* or *my cat*) depends highly *on the ability to mirror the structure of information perceived in the world* [27], and this key indicator can be tuned through the accumulation of new information.

Despite the fundamental differences, classification and identification heavily rely on each other. On the one hand, teleosemantics states that the act of recognising is necessarily prior to the act of classifying [20]. On the other hand, the means employed in identification are often heavily influenced by organised classificational knowledge. For instance, the phrase *'lynxes are large-sized wild cats'* may help someone in correctly recognising a cat-like creature in the forest as a lynx. In this particular example, natural language is used to vehicle classificational knowledge that the receiver can use to improve their identification abilities. Language, for humans, is on par with other perceptual ways of acquiring information, such as vision or hearing [21, 22]. We adopt this point of view for artificial agents in our case study, where we process semi-formal language as a particular form of perceptual input.

3 Classification and Identification in AI

The findings of teleosemantics bear a high relevance to computational models of intelligence. While AI communities have not always been defining the terms *classification* and *identification* in exactly the same manner as above, the respective functions do have AI equivalents. In this section, we map a few important existing AI approaches and tasks to either of these two functions, explaining their differences in the light of teleosemantics. We also show examples of complementary use of identificational and classificational knowledge in existing AI solutions.

In computational systems, classifications are crucial for reasoning, the sharing of knowledge, standardisation, and are generally widely used as vehicles of semantic interoperability, e.g., for data integration. In AI, and in particular in the field of KR, several kinds of representational systems were developed to model classifications as formal and machine-readable grid or tree structures: semantic networks such as in KL-ONE [6] or top-level and domain ontologies (e.g., Dolce [12] or FOAF [7], respectively).

Identification being such a crucial function in processing sensory input in living beings, it is no surprise to find it playing a central role in AI as well. Machine learning (ML) has proven successful for identificational tasks, especially on unstructured 'sensory-like' input such as images or spoken or written natural language [16, 25, 28]. ML classifiers expect such input to be pre-processed ('perceived') as *features* (that map to S in equation 1) and produce *classes* or *clusters* as output (that map to C^I) [5]. ML models, that map to KR^I, are built through the accumulation of input associated to hypotheses ('training'), as foreseen by teleosemantics for identificational representations, instead of the clear-cut classes of classificational KR.

ML is far from being the sole example of identification in AI. *Schema/ontology matching or entity matching*, crucial tasks in practical applications such as data integration, involve identification that maps one or more incoming structures to a set of reference structures. While the inputs of these matching tasks are typically classificational and not perceptual, most matchers analyse them using techniques common for unstructured input, e.g., the extraction of 'features' from ontology labels via NLP [26, 4] and then perform a similarity-based (but not necessarily learning-based) analysis of such features. Note that our teleosemantic model of identification considers the matching of schemas/classes on the one hand and instances on the other hand as essentially the same task over data of different levels of granularity, as opposed to state-of-the-art approaches that regard them as distinct tasks [10]. The need for unifying these tasks has already been recognised in AI in the field of *Structured Machine Learning* [9].

There have been efforts in AI for the mutual reuse of classificational KRs for identificational purposes and vice versa. *Statistical Relational Learning* [13] applies ML to classificational structures. In OntoClean [15], a lot of work has been devoted to defining *identifying* (i.e., *rigid*) *properties* for instances of a certain class (e.g., for an instance of the class *Person*, the *birth date* is an identifying property while *profession* is not as people can change their jobs). In this approach, identificational knowledge is fixed by design as part of classificational knowledge, instead of being derived by gradual accumulation of information. In ontology matching, relying on classificational background knowledge is a common technique for improving precision and recall [10]. Likewise, reusing symbolic knowledge in learning-based (e.g., neural) applications has been a challenging research topic in AI [1, 17, 30, 29]. Giunchiglia and Fumagalli [14] motivate the need for two distinct data layers for the two kinds of knowledge in a context of an ontology built for recognition.

The other direction, namely using identification for building classifications, is manifest in *ontology learning* from unstructured, e.g., textual input [8], ontology matching combined with *repair* [19], and *Inductive Logic Programming* [24]. The latter constructs classificational knowledge by learning from examples, without, however, the use of separate identificational knowledge.

4 The Teleo-KR Architecture

This section aims to formalise the principles of teleosemantics, presented in section 2, as a high-level *Teleo-KR architecture*. We intend the architecture as a frame of reference for interpreting and structuring AI solutions that combine the two essential—classificational and identificational—functions of cognition. The approach is demonstrated in a concrete AI use case in section 5.

Figure 2 shows a high-level schema of the architecture. Rounded boxes correspond to teleological devices, and arrows represent the flow of information. Devices fall into one of three general functional areas or *layers*, modelled within a classic perceptual-cognitive paradigm:

- the *perceptual layer* contains devices that take various forms of input from the outside world: sensory, structured data, unstructured text, etc.;
- the *cognitive layer* with devices that collect and organise information about the world;
- and the *behavioural layer* with devices that act upon the world: moving the agent, communicating with other agents, etc.

The contributions of this paper mostly concern the cognitive layer. As shown in figure 2, the two pivotal devices of teleological representational systems, namely classification and identification, play the role of connecting environmental inputs to behavioural outputs. (Other devices may also be part of this layer, such as one for linguistic reasoning, but they are out of scope for our paper.) Perceptual input first enters into the cognitive layer through the identification device. This design choice encodes the teleosemantic hypothesis that *identification precedes classification* and, more generally, other cognitive and behavioural acts. The fact that in our architecture identification acts as a bridge between perception and other cognitive functions reflects neuroscientific evidence on the complex transition between perception and cognition and is in line with combined perceptual–conceptual theories of knowledge [3].

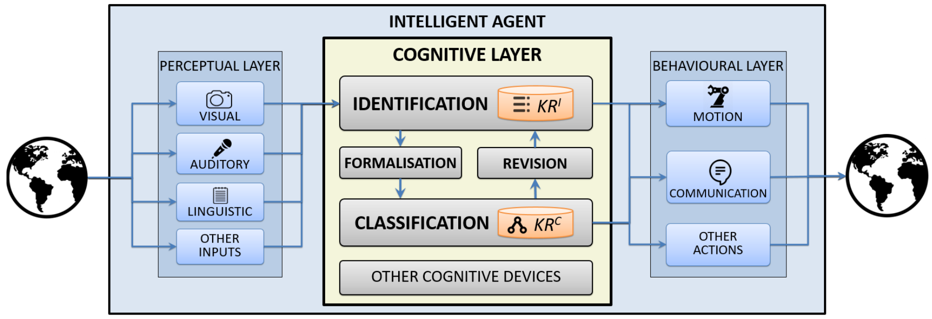

Fig. 1. The Teleo-KR architecture, showing the two pivotal teleosemantic cognitive devices and their relation to perception and behaviour. The schema does not aim an exhaustive description of intelligent agents, hence the inclusion of 'other' devices.

On the other hand, as shown by the architecture, both forms of knowledge play a role in controlling the agent's behaviour. For instance, a communicative act may either be the direct result of instinctive recognition (e.g., shouting upon seeing something frightening) or the vehicle of knowledge in an organised manner.

In this model, classificational knowledge is generated through a process of formalization, which we model as a function $f^F : KR^I \rightarrow KR^C$ *by which classificational knowledge is derived from identificational knowledge.*

Formalisation synthesizes information coming from the external environment, collected through identification during encounters, obtained through various perceptual devices, into a theory about the world. For example, a biologist may observe a living organism from diverse points of view, using an array of sensory inputs (his or her own eyes and hearing, the image provided by a microscope, etc.), before concluding on having discovered an individual of a new species. In the knowledge representation community, this process is known as ontological commitment.

Classifications and deductive thought processes may, in turn, play a role in revising the hypotheses within identificational knowledge, as in the example of the lynx in section 2. Accordingly, we model revision as a function f^R : $KR^C \rightarrow KR^I$ *by which classificational knowledge is used to update identificational knowledge.*

The two processes that interconnect the two forms of reasoning—*formalisation* and *revision*—hide deep open questions about both biological and artificial cognitive systems. Formalisation, i.e., converting a set of incomplete and potentially contradictory hypotheses into a representation of formal classes and relations, amounts to 'making sense' of identification results in a conscious and fine-grained manner. In the context of AI, it is an instance of the *semantic gap problem* that remains only partially solved, especially in the case of deep learning approaches to identification. Likewise, the process of revision, i.e., controlling the inductive process of identification using formally organised rational knowledge, remains ill-understood: one of the major challenges in current AI research is to find efficient ways for plugging in formal knowledge into learning-based systems. These two functions within the Teleo-KR architecture map to an important set of open problems in AI that will remain subject to extensive research in the near future.

5 Case Study

The goals of our case study are: (1) to demonstrate the conceptual power of the Teleo-KR architecture by applying it to a well-known AI task, showing how the latter can be solved through combining classificational and identificational knowledge; and (2) to propose and test a novel idea on the large-scale reuse of existing classificational knowledge for identificational purposes.

The underlying scenario can be described as the *identification of data schemas*: given a set of input *attributes* (or *properties*), find the schema that matches them best. It is a sub-problem of the well-known *schema alignment problem*, used in applications of semantic interoperability, such as data integration or dynamic data matching.

We map this problem onto the Teleo-KR architecture by building a 'teleological AI agent'. We use this agent to simulate a 'cognitive cycle' that starts from perception, identifies the input, builds identificational and classificational knowledge through accumulation and formalisation, respectively, and finally performs a revision of its identificational knowledge to optimise its abilities. We cover the entire cycle through four successive experiments.

Input. As input classificational knowledge we used schemas collected from 15 resources from *Linked Open Vocabularies*[1] (details will be given in each experiment). A major role of such vocabularies, as explained in section 2, is to communicate conventions for interoperability. It thus makes sense to consider them as natural language input received by an intelligent agent through perception, also considering the commitment of teleosemantics on language being on par with other forms of perceptual input (see section 2).

Perceptual preprocessing. We consider the preprocessing of linguistic input as part of perception before identification. Its goal is to generate the *perceptual states* (see section 2 above) that constitute the input of identification. We filtered the input classificational knowledge to retain only (Schema, *attribute*$_1$, …, *attribute*$_k$) relations of labels, e.g., *Person* or *dateOfBirth*. We did not consider attributes inherited from ancestors in order not to bias results by the inheritance hierarchy. Perceiving attribute names as natural language text, we converted them to lowercase, and discarded frequent or meaningless stop words, e.g., $dateOfBirth \rightarrow \{date, birth\}$. The goal was to eliminate surface variations related to orthography, word order, etc. The final output was, for each schema, a bag-of-words vector representation of its corresponding attribute words $Schema_i \rightarrow \left(w_1^{attr}, w_2^{attr}, \ldots, w_n^{attr}\right)$. In machine learning terms,

[1] https://lov.linkeddata.es

Table 1. Accuracies of identificational devices trained (down) and tested (across) on three schema resources.

TRAINING \ TEST	Schema.org	DBpedia	SUMO
Schema.org	*63.77%*	*2.11%*	*1.46%*
DBpedia	*4.19%*	*94.42%*	*7.42%*
SUMO	*1.90%*	*5.82%*	*93,38%*

we consider the words in attributes names as the *features* used by the subsequent identification function. While we could just as well have used a different set of features, optimising this aspect of the setup any further was irrelevant with respect to our experiments.

Identification. We modelled identification essentially as a machine-learning-based supervised document classification task, KR^I being the trained learning model and f^I the learning algorithm. We pre-evaluated multiple algorithms, such as *maximum entropy* or *decision trees*; however, our tests showed that, while the results changed in absolute terms, there was no effect on the overall trends and insights gained. The optimisation of f^I not being of concern to this paper, we finally settled for a decision-tree-based implementation. The training and test sets were all based on the perceived input as described above, with schemas corresponding to output classes and bags of attributes words being the input.

Formalisation. We provide examples to show how acquired identificational knowledge can enrich classificational knowledge. As a form of *ontology learning*, we formalised similarities found among class definitions of various resources by converting them into ontological knowledge of *class equivalence*, *subsumption*, or *semantic similarity*.

Revision. To close the loop, we reused the newly created classificational knowledge to revise identificational knowledge through the optimisation of training data, and thus improve identification results.

Experiment 1: Identification Ability

In our first experiment we trained three identificational devices using three well-known top-level classificational KR resources: *SUMO*[2] (*178 schemas*, *755 attributes*), *Schema.org*[3] (*608*, *877*), and *DBpedia*[4] (*775*, *2861*). We then evaluated each device with respect to their ability to identify schemas, both over themselves (using the same data as for training) and over each other. These evaluations, shown in table 1, quantify the ability of each resource to serve as identificational knowledge.

[2] http://www.adampease.org/OP/
[3] https://schema.org/
[4] https://wiki.dbpedia.org/

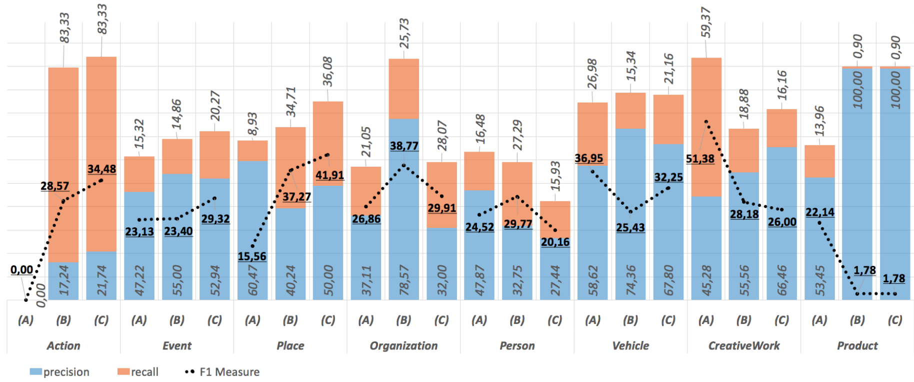

Fig. 2. The effect of accumulation of training data on precision, recall, and F1 for eight core types.

While identification did perform much better, as expected, when the training and test sets were identical, it is also clear that there can be major differences between resources in this respect. Schema.org thus fared much worse for identifying its own schemas. On close analysis, this was due to major overlaps between attribute sets of different schemas, such as the schemas *TVSeries* and *RadioSeries* whose attributes sets were almost identical. The very weak results across resources are, in turn, explained by the relatively low overlap among the schemas and their names (e.g., *Film* in *DBpedia* and *Movie* in *Schema.org* are considered as distinct schemas). This experiment suggests the possibility of a practical tool that evaluates the potential performance of an ontology or a set of schemas in matching tasks. The results may be used, e.g., to finetune schemas in an open-world data integration scenario.

Experiment 2: Knowledge Accumulation

This experiment investigates the effect of *accumulation* of training information on identification results. We increased the size of training sets by merging the three resources from the previous experiment: (A) Schema.org alone; (B) Schema.org + DBpedia; (C) Schema.org + DBpedia + SUMO. We tested the resulting models on a new, more heterogeneous test set consisting of the fusion of 12 vocabularies, some general and some domain-specific, retrieved once again from LOV: *Proton*, *Bibo*, the *Semantic Web for Research Communities*, *SwetoDblp*, the *Comic Book Ontology*, *Linked Earth*, *DNB Metadata Terms*, *Ontology Design Patterns*, *PREMIS*, *EBU*, *Bio*, and *FOAF*. We restricted the evaluation to top-level schemas that were shared by most resources: *Action*, *Event*, *Place*, *Organization*, *Person*, *Vehicle*, *CreativeWork*, and *Product*.

Results can be seen in fig. 2. While accumulation improves the identification of *Action*, *Event*, and *Place*, the improvement is only partial for *Organization* and *Person*, and a deterioration is observed for *Vehicle*, *CreativeWork*, and *Product*. The most salient observation we can make is one well known to the machine

Table 2. Formalisation results: equivalence classes of schemas from Schema.org, derived from identificational similarity scores.

Similar schemas	Similarity
Apartment, SingleFamilyResidence	1.00
Accommodation, House	1.00
Authorize-, Donate-, Give-, Pay-, Return-, TipAction	1.00
Inform-, Invite-, Join-, LeaveAction	1.00
Insert-, Move-, TransferAction	1.00
Comment-, Order-, Reply-, TrackAction	1.00
PropertyValue, QuantitativeValue	0.98
TvSeries, RadioSeries	0.96

learning community: more training data does not systematically lead to higher accuracy. The latter greatly depends on a number of other factors such as input data quality and relevance with respect to the task, how features are defined, the learning algorithm, or the structure of the hypothesis space. In our case, we attribute the low overall scores and the lack of salient improvement of results after accumulation to the high level of heterogeneity of input KRs with respect to the amount of training data.

In conclusion, in a scenario of sparse and heterogeneous identificational knowledge, alternative ways to improve f^I need to be considered beyond the accumulation of more evidence. The Teleo-KR achitecture suggests us the improvement of perception (e.g., through feature engineering) but also the cyclic revision of KR^I using knowledge from KR^C. Our two last experiments illustrate the latter process.

Experiment 3: Knowledge Formalisation

This experiment demonstrates formalisation by reusing the output of identification to enrich classificational knowledge. This operation is analogous to the *ontology repair* or *ontology learning* step that is a regular post-processing feature of many ontology matchers [19].

Table 2 shows sets of schemas from *Schema.org* that were found to be identical or very similar by f^I due to overlapping attributes. The high number of shared attributes found across schemas (the table only shows the tip of the iceberg, as we used a similarity cutoff of 0.95) explains the relatively low identificational power of *Schema.org* in experiment 1. Formalisation converts these observed similarities into acquired classificational knowledge of equivalence, e.g., $TVSeries \equiv RadioSeries$. With a larger-scale analysis that includes property set containment, subsumption relations could also be discovered. Note that in this experiment we only consider *extensional* similarity based on shared attributes. Possible *intensional* similarities and differences could also be taken into account in a more sophisticated formalisation approach that, for example, would consider the semantics of schema names.

Experiment 4: Knowledge Revision

Knowledge revision updates KR^I by classificational knowledge, in our case by the axioms formalised in the previous experiment. We re-trained the *Schema.org*-based model of experiment 1 with schemas found equivalent in step 3. In the training data we replaced each schema with a single one representing their equivalence class, e.g., *TVRadioSeries* or *AccommodationHouse*. We then re-ran evaluations of the retrained *Schema.org* over the original (unmodified) data, and obtained an overall accuracy increase of 2.61%, from 63.77% to 66.38%. With a more aggressive approach to formalisation that does not stop at the similarity threshold of 0.95, accuracy could be increased up to 96.52%. This demonstrates the importance of the formalisation–revision cycle as a means to improve the overall cognitive abilities of the artificial agent.

6 Conclusion and Perspectives

Our paper aimed to reframe a range of tasks and open issues of AI with respect to the functions of *classification* and *identification*. Building on the results of teleosemantics, we defined the two notions, clarified their difference based on analogous functions of natural agents, and demonstrated their pivotal role in AI. Our Teleo-KR architecture proposed a schematic model for AI agents based on the combination of these two functions through *formalisation* and *revision*. We demonstrated the use of the architecture on a set of AI tasks inspired from the well-known problems of *schema identification* and *repair*. The case study also introduced a novel idea for the large-scale reuse of symbolic knowledge resources for identification and learning tasks in general.

Among the potential paths of research opened up by our results, we now present two areas of future work. A first perspective concerns the testing of Rosch's seminal hypotheses on the relatedness of identification with *basic level categories*. We plan to verify her results through computational experimentation using the Teleo-KR framework. A second perspective concerns the notion of reward, a central tool of teleosemantics for the evolution and stabilisation of accumulated knowledge. We wish to formalise reward within the Teleo-KR architecture and investigate parallels with results in AI on reinforcement learning.

References

1. Franz Baader. The description logic handbook: Theory, implementation and applications. *Cambridge university press* (2003).
2. Sebastian Bader and Pascal Hitzler. Dimensions of Neural-Symbolic Integration—*a Structured Survey. arXiv preprint cs/0511042* (2005).
3. Lawrence W. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences* (1999) 22, 577-660.
4. Gábor Bella et al. Language and Domain Aware Lightweight Ontology Matching. *Journal of Web Semantics* 43 (2017): 1-17.
5. Christopher M. Bishop. Pattern recognition and machine learning. *Springer*, 2006.

6. Ronald J. Brachman and James G. Schmolze. An Overview of the KL-ONE Knowledge Representation System. *Readings in artificial intelligence and databases.* Morgan Kaufmann, 1989. 207-230.
7. Dan Brickley and Libby Miller. FOAF vocabulary specification 0.91. (2010).
8. Paul Buitelaar, Philipp Cimiano, and Bernardo Magnini. Ontology Learning from Text: an Overview. *Ontology learning from text: Methods, evaluation and applications* 123 (2005): 3-12.
9. Pedro Domingos. Structured machine learning: Ten problems for the next ten years. *Proc. of the Annual Intl. Conf. on Inductive Logic Programming*, 2007.
10. Jérôme Euzenat and Pavel Shvaiko. *Ontology matching.* Vol. 18. Heidelberg: Springer, 2007.
11. Garrett Evans and John McDowell. *The varieties of reference.* Clarendon Press, 1982.
12. Aldo Gangemi et al. Sweetening Ontologies with DOLCE. *International Conference on Knowledge Engineering and Knowledge Management.* Springer, 2002.
13. Lise Getoor and Ben Taskar. Introduction to statistical relational learning. *MIT Press*, 2007.
14. Fausto Giunchiglia and Mattia Fumagalli. Concepts as (Recognition) Abilities. *FOIS* 2016.
15. Nicola Guarino and Christopher Welty. Evaluating Ontological Decisions with OntoClean. *Communications of the ACM* 45.2 (2002): 61–65.
16. Geoffrey Hinton. Learning multiple layers of representation. *Trends in cognitive sciences* 11.10 (2007): 428-434.
17. Luis da Cunha Lamb. The Grand Challenges and Myths of Neural-Symbolic Computation. *Dagstuhl Seminar 08041 (2008). Recurrent Neural Networks-Models, Capacities, and Applications. Dagstuhl: Schloss Dagstuhl*, 2008.
18. Graham Macdonald and David Papineau et al. *Teleosemantics*, Oxford U. P., 2006.
19. McNeill, Fiona, and Alan Bundy. Dynamic, Automatic, First-Order Ontology Repair by Diagnosis of Failed Plan Execution. *IJSWIS* 3.3 (2007): 1-35.
20. Ruth Garrett Millikan. On Clear and Confused Ideas: An Essay About Substance Concepts. *Cambridge University Press*, 2000.
21. Ruth G. Millikan. Language: a Biological Model. *Oxford Univ. Press*, 2005.
22. Ruth G. Millikan. Learning Language. Published in German translation by Alex Burri as 'Spracherwerb'. *Biosemantik. Sprachphilosophische Aufsätze, Berlin: Surkamp*, 2012. 85-115.
23. Ruth Garrett Millikan. Beyond Concepts: Unicepts, Language, and Natural Information. *Oxford University Press*, 2017.
24. Stephen Muggleton and Luc De Raedt. Inductive logic programming: Theory and methods. *The Journal of Logic Programming* 19 (1994): 629-679.
25. Randall C. O'Reilly. Six principles for biologically based computational models of cortical cognition. *Trends in cognitive sciences*, 2(11), 455-462. 38.
26. Dominique Ritze et al. Linguistic Analysis for Complex Ontology Matching. *CEUR Workshop Proceedings.* Vol. 689. RWTH, 2010.
27. Eleanor Rosch. Principles of categorization. Concepts: *Core Readings 189 (1999).*
28. Linda B. Smith. Learning to recognize objects. *Psychological Science* 14.3 (2003): 244-250.
29. Md Kamruzzaman Sarker et al. Explaining Trained Neural Networks with Semantic Web Technologies: First Steps. *arXiv preprint. arXiv: 1710.04324* (2017).
30. Jennifer Sleeman et al. Entity type recognition for heterogeneous semantic graphs. *AI Magazine 36.1* (2015): 75-86.