

Towards an ontological core for cognitively justified robots

Stefano Borgo¹, Roberta Ferrario¹, Claudio Masolo¹ and Daniele Porello²

¹Laboratory for Applied Ontology, ISTC CNR, Trento

²Department of Philosophy, University of Genova

Abstract

Robots built to interact with humans in everyday life need to organize, manage and elaborate information in ways that are not only reliable but also aligned with humans' understanding of the environment. The paper argues that cognitively motivated ontologies contribute to the development of techniques and to the organization of knowledge bases that take us closer to the construction of robot architectures suitable for smooth interactions with humans.

Keywords

Ontology, Cognitive robotics, DOLCE, Threshold operators

1. Introduction

Increasing attention has recently been paid in robotics towards the development of (everyday, direct) human-robot collaboration. The envisioned future is that cobots (collaborative robots) will help humans in their usual activities, e.g., in goods production as well as housekeeping, and enable smooth social relationships across the human and artificial worlds. It is still unclear how to design cobots for the variety of activities and related possible interactions that are manifested in modern socio-technical systems (STS). This research area is even more challenging when the interaction between robots and humans is not restricted to working scenarios. In everyday situations, human-robot interaction may span the cognitive dimension in its entirety. This short paper presents an approach to cognitive robotics that takes advantage of recent research in applied ontology and knowledge representation (KR).

2. One, No One and One Hundred Thousand

One of the central aims of robotics, especially when considered from the perspective of STS, is to build an object that, when acting, is perceived by humans as a single agent (one), instead of

8th Italian Workshop on Artificial Intelligence and Robotics (AIRO 2021)


✉ stefano.borgo@cnr.it (S. Borgo); roberta.ferrario@cnr.it (R. Ferrario); claudio.masolo@loa.istc.cnr.it (C. Masolo); danielle.porello@unige.it (D. Porello)

🌐 <https://www.istc.cnr.it/people/stefano-borgo> (S. Borgo); <https://www.istc.cnr.it/people/roberta-ferrario> (R. Ferrario); <https://www.istc.cnr.it/people/claudio-masolo> (C. Masolo); <http://www.dif.unige.it/epi/hp/porello/> (D. Porello)

🆔 0000-0002-9175-3096 (R. Ferrario)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

a combination of several interconnected devices and services (one hundred thousand), an agent to interact with rather than something to use (no one). Being observationally coherent and socially predictable is important for artificial embodied agents. A robot has usually a number of components dedicated to collect data about the environment through a variety of sensors and channels (cameras, audio systems, magnetic sensors, wireless communication ...), each with its level of precision, capabilities and limitations (e.g. reflective or transparent objects may be a challenge for optical technology) and possible interferences (from other devices and running processes). Data, rough or pre-processed, coming from different sensors may not be easily merged, e.g., due to differences in tolerance, processing delays or contextual significance of some features. Given that every sensor provides a reading of the environment or fragment of it, for the agent to act coherently there must be a step in which data are reliably integrated and an overall and coherent picture of the environment is generated. Methodologies for data integration, especially when data themselves have to be read from different perspectives of the world (e.g., the physical and the social), are complex. On top of this, methodologies must be rational, robust and produce results that are cognitively acceptable and justifiable. Otherwise the overall artificial system will not be perceivable as a single agent, and what is built is a machine which falls short of being a robot (no one). This means that the process of data integration should follow principles inherited from cognitive sciences. Data integration is obviously not enough to reach an interpretation of the available information enabling a meaningful interaction with humans. In order to reach the latter, the robot's knowledge should be enriched with input concerning the socio-cultural environment in which it operates, including such diverse information as attribution of roles, working practices, social rituals, taboos etc. Given the heterogeneity of the information that robots should use when interacting with humans, a methodology for organizing, deploying, and representing it is mandatory. In the next sections, we will propose a cognitively motivated ontology as an appropriate means to tackle such complex issue.

3. Cognitive Ontology in the Agent's Information Flow

Techniques developed in applied ontology and KR are already exploited in robotics to some extent [1]. One of the purposes of applied ontology is the transparent and robust organization and classification of data coherently with the conceptualization the observer has of reality [2]. This effort led to develop several ontology-based methodologies which, paired with techniques for decision making, can improve sensor data classification, data analysis and data integration.

Generally speaking, a formal ontology is a general, reliable and well organized conceptual system. This means that the ontology includes the most usable, domain-independent and widely applicable concepts (generality), is expressed as a logical theory with Tarskian-style semantics, has a rich axiomatization and carefully analyzed formal consequences (reliability); and is constructed following explicitly motivated philosophical principles (justified organization). Overall, an ontology relies on the symbolic representation of data, and its aim is to formalize a view of the world, i.e., a way to coherently understand reality and what one expects to be possible. The ontology DOLCE [3] is an example of such systems. A knowledge-base built from this ontology can model information from the agent's perspective.

As said, formal ontologies are written in logical languages such as first-order logic or, to

ensure the tractability of the reasoning services, in weaker fragments of first-order logic. A fundamental family of languages for ontologies is that of Description Logics (DL). E.g., the Web Ontology Language (OWL 2) [4, 5] corresponds essentially to the $SR\mathcal{OIQ}^{(D)}$ DL [6].

Logic-based ontologies are capable of specifying hard constraints on individuals, e.g. the rigid distinction between an object and its qualities, or an object and its components. To cope with concepts learnt from examples or data, one can introduce dedicated operators like the *Weighted Threshold Operators*. Informally, these operators take a list of relevant concepts in the ontology and treat them as features to define the newly learnt concept.

In [7], we started studying weighted logics to provide cognitively meaningful representations of concepts in ontologies, while in [8] and [9], we studied threshold operators in the context of DLs. In brief, if $C_1 \dots C_n$ are concept expressions, $w_1 \dots w_n \in \mathbb{R}$ are weights, and $t \in \mathbb{R}$ is a threshold, the new concept $\mathbb{W}^t(C_1 : w_1, \dots, C_n : w_n)$ (called a ‘Tooth expression’) classifies the individuals d such that $\sum\{w_i : C_i \text{ applies to } d\} \geq t$. As a toy example, assume we wish to capture a common-sense definition of *mug*. We list a number of relevant features with the associated weights to indicate their relevance for an object to be classified as a mug. This is expressed by the TBox axiom:

$$\mathbf{Mug} \equiv \mathbb{W}^2(\exists \text{contains.}\mathbf{Coffee} \sqcup \mathbf{Tea} : 1, \\ \exists \text{hasBase.}\mathbf{Circular} : 1, \exists \text{hasPart.}\mathbf{Handle} : 1, \exists \text{isLargerThan.}\mathbf{Cup} : 1) \quad (1)$$

That is, a mug is associated with features like: it contains coffee or tea, it has a circular base, it has a handle, it is larger than some cups. The definition adds that, for being a mug, any two of these four features suffice (here, 1 is the weight assigned to each feature, and 2 is the threshold).

Threshold operators have been extensively studied in the context of circuit complexity theory (e.g. [10]), and they are also known in the neural network community by the name of *perceptrons* (cf. e.g. [11]). Extensions of DLs with threshold operators have been discussed also in [12, 13]. Classifications via the threshold concepts can be construed as *knowledge-dependent*, cf. [9]: the values depend on the knowledge base, so the classification depends on the knowledge available to the agent. This aspect highlights the situated *contextual* nature of classification tasks.

The benefits of Tooth expressions are briefly summarized as follows: *i*) they provide a way to compactly define concepts that are human accessible: instead of writing a possibly long disjunctive normal form (DNF) of concepts, we list the weighted relevant features and a threshold; *ii*) definitions provided in terms of Tooth can be grounded on cognitive theories of concepts and categorization, like the theory of prototypes or the theory of exemplars, cf. [14]; *iii*) adding Tooth expressions to rich DLs does not raise the computational complexity of the reasoning services (cf. [9]); *iv*) the weights of the Tooth expressions can be learnt from examples or past experience of an agent, so Tooth expressions can bridge symbolic knowledge with experiential data. For details, we refer to [9].

4. Ontologizing the robot’s architecture

Even if traditional problems in robotics (navigation, object detection, obstacle avoidance, object grasping etc.) were completely solved, the types of information that a cobot should be able to manage remain impressive [15]. Since our goal in this paper is to suggest an ontology-based approach to information in robotics, and to exemplify it in the context of information integration,

we now show how this can be implemented in a generic robot architecture (Fig. 1) . It should be clear from the earlier sections that our work is essentially about the meaning of information (with impacts on its collection, organization and manipulation), which is typically a concern in the robot’s modules controlled by knowledge representation techniques [16].

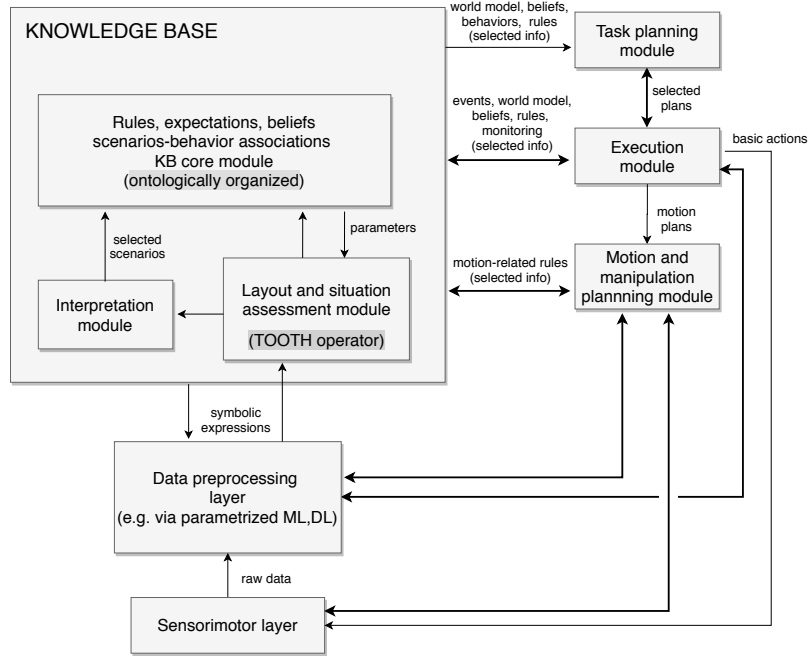


Figure 1: Information flow in a general-purpose robot architecture (adapted from [15]).

In Fig. 1 (bottom) data is collected by a sensorimotor module and then preprocessed by a parametric module, typically relying on subsymbolic approaches. The information produced by this latter module is available in logical format and is essentially about geometric information and sensor-dependent features. This information is processed by the layout and situation assessment module which identified objects and the overall layout of the (detectable) environment. The Tooth expressions described in the previous sections are devised to integrate information at this step. The layout and situation description is then passed to the interpretation module and to the core of the knowledge base (which we do not discuss). The role of the interpretation module is to elicit the interpretation of the situation, what we call a *scenario*. Briefly put, a situation is the mere list of objects and their places in an environment (there are people forming a line), the scenario is the social interpretation of what is going on (there is a queue), allowing to associate to the actual state the relevant rules, behaviors and expectations as described in [15]. The rest of the architecture, which we do not discuss here, describes the connections across the task and geometric planners and the execution module.

In this example, the discussion of section 3 refers to a module where both classical DL expressions and Tooth expressions are used. While the DL expressions represent factual knowledge (the *ABox*) or hard laws that regulate the universe of discourse, the Tooth expressions define concepts under which entities can be classified on the basis of the data (*ABox* statements)

returned by the data preprocessing layer.

First note that, since Tooth expressions are legitimate logical concept constructors, they can be exploited by, and can exploit, the reasoning capabilities of the main knowledge base module. For instance, given the definition of **Mug** in equation (1), from the ontological knowledge present in the KB core module one can infer that mugs are physical objects, that they are located in space, that they have container functions, etc. Moreover, the ontological knowledge can be used to deductively close the *ABox* returned by the preprocessing layer empowering the classification mechanism. E.g., one can deploy known correlations among qualities to infer objects' features that the sensors were unable to directly observe.

Second, Tooth definitions of concepts can be automatically built by running supervised learning algorithms on the dataset of *ABox* statements provided by the preprocessing layer (at a given time) with the major advantage of being in principle human readable and explainable. Moreover, knowledge about planned actions and expectations can be used to modify or refine these definitions (for example acting on the weights used by the Tooth expressions) on the basis of the success of the actions undertaken by the robot. Finally, the Tooth definitions can be used to guide the robot to focus on the acquisition of specific information with the goal of disambiguating doubtful yet critical classifications. For instance, suppose that the robot is unable to classify an object *a* because no threshold has been reached. Starting from the concepts with the highest classification degree for *a* and possibly from knowledge of what is typically present in a scenario, the robot can plan to make more observations and even run a check of its sensors for possible malfunctioning.

Third, even the robot architecture could be refined to match assumptions largely adopted by cognitive theories of concepts and categorization. For instance, cognitive theories usually make a distinction between *concepts* (e.g., dog, mug, tree), *attributes* (e.g., color, shape, texture), and *attribute-values* (e.g., crimson, round, smooth). The classification of an object under a concept is determined on the basis of its attribute-values and of the relevance and typicality of these values. Attributes and attribute-values are strongly linked to the human perceptive system, while concepts are cognitively more complex. This contrast can be simulated by assuming that (i) by being closely linked to the available sensors, the preprocessing layer returns assertions based on DL-concepts representing attribute-values, while (ii) the complex concepts are defined via Tooth-expressions involving only the DL-concepts corresponding to attribute-values with weights representing the relevance and typicality for that concept.

While the approach we have presented is cognitively interesting and adds the flexibility of a parametrized classification, it also raises a series of problems. The elaboration of the output of the sensors performed at the preprocessing layer could result in a set of assertions that are inconsistent with the general knowledge in the KB core module or that generate inconsistent classifications under Tooth-defined concepts. This happens, for instance, when an individual satisfies enough features associated with concepts that are deemed disjoint by an axiom in the KB core module. There can be several heterogeneous causes that lead to this kind of conflict: the disjointness axiom might be wrong, the tagging of the data might be mistaken, the Tooth definition of the concepts might be incomplete, the scenario might have changed etc. There are approaches to overcome these conflicts, for instance by applying *judgment aggregation* techniques in decision theory [17] that may use meta-information about the sensors, for instance, their reliability in given contexts (see [18]).

References

- [1] A. Olivares-Alarcos, D. Beßler, A. Khamis, P. Goncalves, et. al., A review and comparison of ontology-based approaches to robot autonomy, *Knowledge Engineering Review* 34 (2019).
- [2] N. Guarino, Formal ontology in information systems, in: N. Guarino (Ed.), *Proceedings of the Second International Conference on Formal Ontology in Information Systems*, IOS Press, 1998, pp. 3–15.
- [3] S. Borgo, C. Masolo, Foundational Choices in DOLCE, in: S. Staab, R. Studer (Eds.), *Handbook on Ontologies*, 2nd ed., Springer Verlag, 2009, pp. 361–381.
- [4] O. W. Group, et al., *OWL 2 Web Ontology Language Document Overview: W3C Recommendation 27 October 2009* (2009).
- [5] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, S. Rudolph, *OWL 2 web ontology language primer*, W3C recommendation 27 (2009) 123.
- [6] I. Horrocks, P. F. Patel-Schneider, Reducing OWL entailment to description logic satisfiability, in: *International semantic web conference*, Springer, 2003, pp. 17–29.
- [7] C. Masolo, D. Porello, Representing concepts by weighted formulas, in: S. Borgo, P. Hitzler, O. Kutz (Eds.), *Formal Ontology in Information Systems - Proceedings of the 10th International Conference, FOIS 2018, Cape Town, South Africa, 19-21 September 2018*, volume 306 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2018, pp. 55–68.
- [8] D. Porello, O. Kutz, G. Righetti, N. Troquard, P. Galliani, C. Masolo, A toothful of concepts: Towards a theory of weighted concept combination, in: *Proceedings of the 32nd International Workshop on Description Logics*, volume 2373, CEUR-WS, 2019. URL: <http://ceur-ws.org/Vol-2373/paper-24.pdf>.
- [9] P. Galliani, O. Kutz, D. Porello, G. Righetti, N. Troquard, On knowledge dependence in weighted description logic, in: *Proc. of the 5th Global Conference on Artificial Intelligence (GCAI 2019)*, 2019, pp. 17–19.
- [10] H. Vollmer, *Introduction to circuit complexity: a uniform approach*, Springer Science & Business Media, 2013.
- [11] C. M. Bishop, *Pattern recognition and machine learning*, Springer Science+ Business Media, 2006.
- [12] F. Baader, G. Brewka, O. F. Gil, Adding threshold concepts to the description logic \mathcal{EL} , in: C. Lutz, S. Ranise (Eds.), *Frontiers of Combining Systems*, Springer International Publishing, Cham, 2015, pp. 33–48.
- [13] F. Baader, A. Ecke, Reasoning with prototypes in the description logic \mathcal{ALC} using weighted tree automata, in: A.-H. Dediu, J. Janoušek, C. Martín-Vide, B. Truthe (Eds.), *Language and Automata Theory and Applications*, Springer International Publishing, Cham, 2016, pp. 63–75.
- [14] G. Righetti, D. Porello, O. Kutz, N. Troquard, C. Masolo, Pink panthers and toothless tigers: Three problems in classification, in: *Proc. of the 5th International Workshop on Artificial Intelligence and Cognition*, Manchester, September 10–11, 2019.
- [15] S. Borgo, E. Blanzieri, Trait-based module for culturally-competent robots, *International Journal of Humanoid Robotics* 16 (2019) 1950028.
- [16] K. Rajan, A. Saffiotti, Towards a science of integrated ai and robotics, *Artificial Intelligence*

247 (2017) 1 – 9.

- [17] D. Porello, N. Troquard, R. Peñaloza, R. Confalonieri, P. Galliani, O. Kutz, Two approaches to ontology aggregation based on axiom weakening, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden., 2018*, pp. 1942–1948. URL: <https://doi.org/10.24963/ijcai.2018/268>. doi:10.24963/ijcai.2018/268.
- [18] C. Masolo, A. B. Benevides, D. Porello, The interplay between models and observations, *Applied Ontology* 13 (2018) 41–71.