CEUR-WS.org/Vol-2708/robontics8.pdf

Robot Meets World[1]

Jona THAI[a], Michael GRÜNINGER[a]
[a] *Department of Mechanical and Industrial Engineering, University of Toronto, Ontario, Canada M5S 3G8*

Abstract. The representation of the world in which an autonomous robot interacts with other objects within the world requires the integration of perception, reasoning, and action. The challenge of designing a suite of ontologies that are expressive enough to support perception and reasoning is a preeminent challenge for the knowledge representation and applied ontology communities. In this paper, we propose a series of challenge problems oriented around the ways in which a robot encounters the physical world. We outline an approach to the design and application of ontologies that can be used to represent the problems and their solution.

Keywords. ontology process mereotopology robotics first-order logic PSL

1. Introduction

Imagine waking up in a well-lit, unrecognizable room. Although terrifying and stress-inducing, a quick scan should allow you to quickly devise a plan of action. You are able to discern what sharp objects to avoid, what objects are movable, ”free space” and what would lead to a ”free fall”. Much of these common-sense reasoning is accumulated through trial-and-error over the years. However, even more of this common-sense is taught to us – for instance, we don’t need to get stabbed by a knife ourselves to understand that it is fatal. A small portion of our intelligence is genetic – for instance, we figure out how to walk and grab things without explicit teaching or extensive experience. This combination of spatial, temporal and physical knowledge allows us to reason about the properties of spaces and objects that we have never encountered. The question is, what is the minimum(albeit non-exhaustive) knowledge needed to represent a robot as an embodied cognitive agent, one that possesses a physical form, is aware of it, and interacts with other physical objects and other embodied cognitive agents? To this end, we will explore several scenarios and propose relevant ontologies.

2. Informal Specification of Reasoning Problems for Robotics

Assuming a robot has successfully replicated our spatiotemporal and physical reasoning capability, what is the minimum we can expect it to be able to do? In this section, we present an initial set of reasoning problems that can be used to determine the suite of ontologies that we will need.

---

[1] Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Problem 1 *Path Planning*
Given a room containing a set of objects, find a path that the robot can take from one location to another which avoids contact with any objects.

Problem 2 *Relating Self to Object*
Given an Article A, and a room containing a set of objects, find a path the robot can take while carrying Article A from one location to another, while avoiding contact with any objects.

Problem 3 *Up, Down and All Around*
Given a tall Obstacle A, wide Obstacle B and ledge C, find the optimal path the robot can take.

Problem 4 *Bob The Builder*
Given a set of instructions and a set of objects, the robot should build a composite object.

Problem 5 *Not Whole But Hole*
Given a room with holes, find the optimal path the robot can take to the other side of the room.

Problem 6 *Tool Building*
Given a room with a set of objects, and a chasm, assemble the objects in a way that the robot can cross the chasm without falling into it.

Problem 7 *Path Planning 2*
Given a room with a set of moving objects, find a path the robot can take from one location to another which avoids contact with any objects.

Problem 8 *Breaking Barriers*
Given a room with a set of objects forming a wall, move the objects or the robot in such a way that it can traverse to the other side.

Problem 9 *Lock(e) and Key*
Given a Lock e, and a two vastly different keys, the robot must choose the right key the first time.

Problem 10 *Weigh the Way*
Given a chasm and several objects of varying length and weight capacity, the robot needs to accurately pick objects that will aid it in successfully crossing the chasm.

Problem 11 *Together is Better*
Given a wall W that is too tall for a single robot to traverse without assistance, find a way for two robots to traverse the wall.

We believe that these problems address the minimum ontological commitments for a robot to replicate our ability to reason about space and movement. The next step is to design the weakest ontology possible for the robot to reason about solutions to the above set of problems.

3. Existing/Related Work

Organizing knowledge around robotics is not a novel pursuit. Below we explore three largely known ontologies within the field of robotics.

3.1. Core Ontology for Robotics and Automation(CORA)

CORA is an ontology specifically focused on the fields of industrial robotics, service robotics, and autonomous robotics. It formally specifies mechanical components of robots and robotics, types of robots and types of robot interactions. This is ideal for classifying different types of robotic designs and movements. This is achieved through an alignment with upper ontologies such as SUMO, and OWL. [1]

3.2. KnowRob(2)

KnowRob is notable for its contributions in providing a heterogenous approach to semantics in robotics. It utilizes a trove of ”narrative enabled episodic memories”, or NEEM, for robots to learn from. This data is collected from multiple robotic systems, and stored as RDF triples which are queried over the OpenEASE platform. Although effective, this combination of software and ontologies is difficult to re-use and replicate, due to barriers of entry of programming language and platform. [2]

3.3. OpenRobots Common Sense Ontology(ORO)

The OpenRobots Common Sense Ontology is closely related to KnowRob, both in sharing of concepts and design concept. However, it differs from KnowRob in that it focuses on human-robot interaction. It is aligned with the now defunct OpenCyc upper ontology. It achieves its purpose by developing two main ontologiesn- one focusing on the theory of the mind and one of ”common-sense”. The former helps develop context and make sense of ambiguity in instructions, whereas the latter promotes an ability to reason about actions to take. As such, ORO’s goal differs slightly from our proposed problem, as we focus on problems in terms of robot-physical world interaction. [3]

3.4. What Ontologies Do We Need

Given the combined capability of these ontologies, one may wonder why we are proposing yet another ontology for robotics. However, it should be noted that in terms of application, we differ from ORO, in the sense that we focus on robot-physical world interactions as compared to robot-human interaction. Moreover, we differ from CORA in the sense that we do not seek to create a robotics ontology, but rather ontologies for robots (that will support common-sense reasoning). In addition, all of the above ontologies do not contain extensive first-order axiomatizations, which limits expressivity. KnowRob cicumvents this through their hybrid, software-oriented approach. Although effective, system complexity and different programming languages prove a barrier in reusability. As such, our design of a reasoning framework purely based on first-order logic is novel.

How can we be sure that our set of designed ontologies are right for the application? Following [4], an ontology works well if it does not generate unwanted models, and does not constrict upon desired models. For example, mainstream designs of spatial reasoning often utilize Euclidean 3D geometry. This intuition is not unfounded, as Newtonian mechanics occurs within this framework. However, this model is too strong, in the sense that we do not need such formal geometry to represent overlapping regions. The same result can be achieved with a mereotopology. At the same time, Euclidean geometry alone is too weak for human-like spatial reasoning. This is because it does not contain axioms that can differentiate between physical objects and the regions they occupy.

4. PSL Ontology

The PSL Ontology [5] is a modular first-order ontology that axiomatizes a set of intuitive semantic primitives that is adequate for describing the fundamental concepts of processes. Two of the theories within the PSL Ontology play key roles in this paper – $T_{occtree}$ (theory of occurrence trees) and T_{disc_state} (theory of fluents and discrete change) [2].

Within the PSL Ontology, an occurrence tree Γ is a partially ordered set of activity occurrences, such that for a given set of activities, all discrete sequences of their occurrences are branches of the tree. An occurrence tree contains all occurrences of all atomic activities; activity occurrences that are elements of an occurrence tree are called arboreal activity occurrences.

Most applications of process ontologies are used to represent dynamic behaviour in the world so that intelligent agents may make predictions about the future and explanations about the past. In particular, these predictions and explanations are often concerned with the state of the world and how that state changes. The PSL core theory T_{disc_state} is intended to capture the basic intuitions about states and their relationship to activities.

Within the PSL Ontology, the notion of state is represented by reified fluents. Intuitively, a change in state is captured by fluents that are either achieved or falsified by an activity occurrence. The *prior* relation is used to specify the fluents that are intuitively true prior to an activity occurrence and the *holds* relation specifies the fluents that are intuitively true after an activity occurrence.

Furthermore, a fluent can only be changed by the occurrence of activities. Thus, if some fluent holds after an activity occurrence, but after an activity occurrence later along the branch it is false, then an activity must occur at some point between that changes the fluent. This also leads to the requirement that the fluents holding after an activity occurrence will be the same fluents that are prior to any successor occurrence, since there cannot be an activity occurring between them.

The key step in applying the PSL Ontology to the problems in Section 2 uses the approach presented in [6], which begins with an ontology about a specific domain, and then characterizes the ways in which activities can possibly change relationships among objects in the domain. By identifying models of the domain ontology with states within the process ontology, we can use properties of models to classify activities.

Axioms in the ontology are mapped to state constraints, which are universal sentences containing only *prior* literals and a unique activity occurrence variable. For any domain ontology T, there exists a set of state constraints known as the domain state ontology. The domain process ontology is designed by focussing on the relationship be-

---

[2] This paper uses the axiomatization of the PSL Ontology found at `http://www.mel.nist.gov/psl/psl-ontology/`.

tween models of the domain ontology and the states associated with activity occurrences in the occurrence tree.

5. Formal Specification of reasoning problems

We can now provide a formal specification of the reasoning problems from Section 2, and identify the ontologies that are required for the specification.

5.1. Revisiting Problem 1

The first step in the formalization of Problem 1 is to identify the basic ontological commitments that are required to specify the problem.

We begin with the existence of a room that contains a set of objects. As observed by [7], this leads to the distinction between physical objects and the spatial regions that they occupy, and hence raises the approach of mereotopological pluralism – there is a mereotopology on physical bodies and another distinct mereotopology on spatial regions. Moreover, the notion of location requires a mereotopological harmony principle, that is, the existence of a homomorphism from the mereotopology of physical objects to the mereotopology of spatial regions. As a result of these requirements, we need to reuse the T_{occupy} ontology, which the has signature
$\{\textbf{region},\textbf{part},\textbf{physical\_body},\textbf{physical\_part},C,\mathit{physical\_C},\mathit{occupies}\}$

The next aspect of Problem 1 focusses on the statement "that the robot can take from one location to another". In other words, we need an ontology that can represent how objects can change location, that is, an object can occupy one region at one point in time, and then occupy a different location at another point in time. The design of an ontology that satisfies this requirement is based on the methodology of [6] and [8] by which a domain ontology (in our current case being T_{occupy}) is used to specify a domain state ontology. In particular, axioms of the Occupy State Ontology include state constraints such as

$$(\forall x,y,r_1,r_2,o)\, prior(tphysical\_part(x,y),o) \wedge prior(loc(x,r_1),o) \wedge prior(loc(y,r_2),o) \supset prior(tpart(r_1,r_2),o) \tag{1}$$

$$(\forall x,y,o)\, prior(contact(x,y),o) \equiv (\exists r_1,r_2)\, prior(loc(x,r_1),o) \wedge prior(loc(y,r_2),o) \wedge prior(tC(r_1,r_2),o) \tag{2}$$

The initial for the problem is specified as:

$$prior(loc(Robot,R1),Occ_{initial})prior(loc(Chair,R3),Occ_{initial}) \tag{3}$$

while the goal state is:

$$(\exists o)\, holds(loc(Robot,R4),o) \tag{4}$$

The Motion Domain Process Ontology classifies activities by the ways in which objects can change their location:

$$(\forall a)motion(a) \equiv ((\forall o)occurrence\_of(o,a) \supset (\exists x,r_1,r_2)falsifies(o,loc(x,r_1)) \wedge achieves(o,loc(x,r_2)) \tag{5}$$

The spatiotemporal objects that represent the path that the robot takes is axiomatized as a class of complex activities within the PSL Ontology:

$$(\forall a)\, path(a) \equiv ((\forall a')\, subactivity(a',a) \wedge primitive(a') \supset motion(a')) \tag{6}$$

Finally, the constraint in Problem 1 is axiomatized by:

$$(\forall x,y,o,o_1,a)\, path(a) \land occurrence\_of(o,a) \land subactivity\_occurrence(o_1,o) \supset \neg holds(contact(x,y),o) \quad (7)$$

5.2. Revisiting Problem 2

To identify which ontologies are needed to satisfy Problem 2, we can first attempt to break down the question as per Problem 1. Again, we begin with a room containing a set of objects, except with the addition of Article A. The next aspect of the problem then explains the role of Article A - the robot has to get across the room, without bumping any of the obstacles, while carrying Article A.

Clearly, this problem is an extension of Problem 1. In addition to understanding the relation between an object and the physical space it occupies, the robot needs to now understand the relation between itself and another object. Our approach to this problem utilizes the mereotopology of connected objects proposed in [9], in which objects are connected iff they can be summed together to form a new object. The idea is to consider the robot and the object it is holding to constitute a new object that is the sum of the robot and the object. In this way, a mereotopology is sufficient to represent the fundamental ontological commitments for a robot to be holding an object.

5.3. 1, 2, 3...Infinity!

Perhaps evident from the previous two problems, the set of 12 competency questions above build upon each other. Due to time and space constraints, we will not explore the remaining questions in detail, but offer suggestions as to how these questions may be viewed and tackled. For instance, Problem 5 and 9 require explicit formalization of non-physical obstacles e.g. holes. To comprehend a hole, the robot needs to understand the distinction between material and immaterial objects, perhaps through algebraic topology. Problem 9 requires an additional ontology for shape to map complementary relations between concavity and convexity. Another example is Problem 10 and 11 - both of these problems require the robot's self-awareness regarding its own capabilities relative to the obstacles before it. In addition, Problem 11 introduces the concept of collaboration between robots, an ability of definite value in large scale problems e.g. construction.

6. Conclusion

A typical approach to the application of ontologies to the representation and solution of commonsense reasoning problems is to use the strongest possible ontology (such as full Euclidean geometry or ordered real closed fields) to represent the world in which a robot is situated. In this paper we have taken a different perspective and sought the weakest possible ontology that is sufficient to axiomatize the intended semantics presupposed by the natural language statements of challenge problems for robotic knowledge representation and reasoning. In particular, we have seen how mereotopology, location, and process ontologies are sufficient to axiomatize the fundamental ontological commitments of robot path planning and limited interactions of robots with physical objects in their environment. In the spirit of the Physical Turing Test [10], we have proposed a series of challenge problems that will serve as a platform for the design and evaluation of ontologies to support the representation and solution of these problems.

References

[1] Prestes E, Carbonera JL, Rama Fiorini S, M Jorge VA, Abel M, Madhavan R, et al. Towards a core ontology for robotics and automation. Robotics and Autonomous Systems. 2013;61(11):1193 – 1204. Ubiquitous Robotics. Available from: `http://www.sciencedirect.com/science/article/pii/S0921889013000596`.

[2] Beetz M, Beßler D, Haidu A, Pomarlan M, Bozcuoglu AK, Bartels G. Know Rob 2.0 — A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents. 2018 IEEE International Conference on Robotics and Automation (ICRA). 2018:512–519.

[3] Lemaignan S, Ros R, Mösenlechner L, Alami R, Beetz M. ORO, a knowledge management platform for cognitive architectures in robotics; 2010. p. 3548–3553.

[4] Aameri B, Gruninger M. Which Ontology is Right for Me? In: AAAI; 2015. .

[5] Gruninger M. Ontology of the Process Specification Language. In: Handbook on Ontologies in Information Systems. Springer-Verlag; 2003. .

[6] Aameri B. Using Partial Automorphisms to Design Process Ontologies. In: Proceedings of the 7th International Conference on Formal Ontology in Information Systems. IOS Press; 2012. p. 309–322.

[7] Aameri B, Grüninger M. Location Ontologies based on Mereotopological Pluralism. Applied Ontology. 2020;15(1):1–50.

[8] Aameri B, Gruninger M. A First-Order Axiomatization of Change in Mereotopology. In: Proceedings of the 27th International Workshop on Qualitative Reasoning; 2013. p. 115–121.

[9] Gruninger M, Chui C, Ru Y, Thai J. A Mereology of Connected Structures. In: FOIS; 2020. .

[10] Ortiz CL. Why We Need a Physically Embodied Turing Test and What It Might Look Like. AI Magazine. 2016;37(1):55–62.