Heterogeneous ontologies and hybrid reasoning for service robotics: the EASE framework

John Bateman, Michael Beetz, Daniel Beßler,
Asil Kaan Bozcuoğlu, and Mihai Pomarlan

Everyday Activities Science and Engineering (EASE),
Collaborative Research Centre, University of Bremen,
D-28359 Bremen, Germany
WWW: http://ease-crc.org/

Abstract. As robots are expected to accomplish human-level manipulation tasks, the demand for formal knowledge representation techniques and reasoning for robots increases dramatically. In this paper we describe how to make use of heterogeneous ontologies in service robotics. To illustrate the vision, we take the action of *pouring* as an example.

Keywords: ontology, robotics, heterogeneity, hybrid reasoning

1 Introduction

In this paper we address the challenge of providing robotic agents with the substantial and diverse kinds of knowledge required for the execution of intelligent activities. Whereas ontologies are regularly employed for aspects of such knowledge, many further distinct sources of information equally important for intelligent action have not so far been considered from an explicitly ontological perspective. The framework under development in our Everyday Activities Science and Engineering research centre (EASE) attempts to move in precisely this direction in order both to extend what is typically understood as the scope of 'ontology' and to increase the degree of sophistication and naturalness of robotic actions and interactions with humans and other robots.

The distinct kinds of knowledge relevant can be characterized succinctly with the help of an example. Consider an instruction to a service robot to 'pour' some substance into a container – as in 'pour the coffee into the mug'. In order to successfully and appropriately perform this action, the robotic agent must first have knowledge concerning 'coffee' and 'mug' and mechanisms for grounding these linguistic phrases into the practical situation to deliver references, as well as knowledge that 'pouring' is a movement involving the entity poured and several other entities not explicitly referred to in the linguistic utterance, such as the container from which and with which the pouring takes place. Abstract knowledge of this kind is often provided in ontology representations.

In addition, however, the robot also needs to draw on substantial *experiential* knowledge: for example, if the container holding the substance to be poured is held too high, there is likely to be spillage – knowledge of this kind can be derived both by practical experimentation in the real world and by simulation involving naive physics. Moreover, the particular fine-grained actions to be performed vary depending on the specific robot involved: grasping and movement instructions at this level also need to be provided in order to more finely articulate the abstract, or general, notion of pouring. The actual pouring action performed depends further on the kind of material being poured (coffee acts differently to pancake mixture or beaten eggs), on the destination of the pouring (into a mug, onto a hotplate for frying, into a further liquid while being mixed), etc. The knowledge of differences in the pouring action here is similarly drawn from combinations of abstract knowledge, practical experiential knowledge and simulation of what could occur given the physical parameters of the starting situation.

Most of these kinds of knowledge are not currently considered formally as components of the ontological characterizations provided, which raises issues concerning how these sources of information can be gathered and maintained as a system gains experience and how their interaction can be specified in order to flexibly solve problems. We argue that presenting all such knowledge from the perspective of linked ontological specifications of distinct kinds offers a powerful, elegant, and more integrative framework for robotic agent design.

2 Conceptual Apparatus

Several terminologies and concepts are necessary to enable robotic applications to make use of formal knowledge representation and reasoning techniques.

2.1 Ontologies: purposes and organization

Intelligent systems require rich sources of knowledge in order to reason about solutions to the problems they are designed to deal with. Such knowledge can become both extensive and complex in its internal organization, requiring effective software engineering techniques for development, maintenance and use. In *Artificial Intelligence* (AI), it has long been common to employ notions of formal ontology to meet these challenges [12]. In the classic definition of Gruber, an ontology is an "explicit specification of a conceptualization" [11], that formalizes the types, properties, and interrelationships of concepts for particular domains and which support automated reasoning techniques for problem solving within those domains.

Ontologies share many structural similarities despite being written in different formats and languages. Most ontologies describe instances, classes, attributes, and relations. Instances include objects such as humans, mugs and coffee as well as abstract individuals like numbers. Classes are used to organize objects and classes. For instance, an abstract person or a specific person called Max Mustermann may be an instance of class *Person*. Again, a class called *Mug* may have a super-class called *Container*. Instances and classes are generally described by relating them to other parts and conceptual properties. These relations are often called attributes, although they may be independent instances. Relationships between instances specify how instances are related to each other.

A concrete example is shown in Figure 1 in which a *Pouring* class is illustrated with its related superclasses in the large-scale Cyc ontology [17].

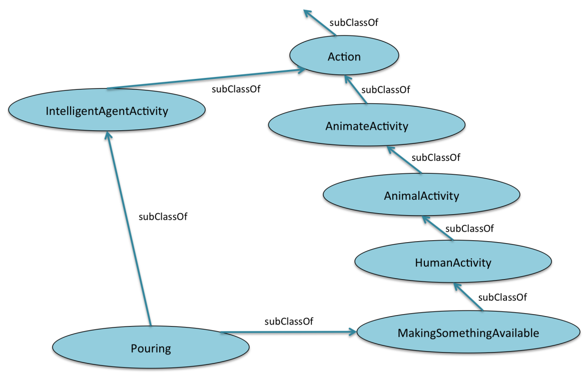

Fig. 1. The 'pouring' action and super-classes in the Cyc ontology [17].

Ontologies are generally employed for several distinct tasks in addition to capturing the relationships and entities within some domain. Some ontologies are explicitly designed as *top level* ontologies: they provide broad organizational constraints for modeling decisions independently of specific domains, thereby imposing a well motivated modeling discipline for any particular concepts, relations or combinations of concepts and relations required. Top level ontologies explicitly orientate towards philosophical discussions concerning more traditional ontological topics. The most well established top-level ontologies are Guarino, Masolo and Borgo's DOLCE [6], Smith and colleagues' BFO [1], and Pease's SUMO[20]. The provision of top level ontologies is intended to increase *interoperability* across distinct system components, one of the primary reasons put forward for adopting ontologies in information systems at all. When separate components are specified in a manner consistent with the modeling decisions of a top level ontology, the likelihood that information may be combined without complications is increased and alignment across ontologies is more likely to be effective [8].

As ontologies are developed, they can exhibit considerable size and complexity. Maintaining them is then a challenge in its own right and a broad variety of tools and, to a lesser extent, techniques to support this task have emerged. Many ontologies are now organized into distinct *modules* which are then combined to cover domains or areas of concern. Modules generally reflect 'theories' of some particular subdomain and may themselves be more or less general in application – the Cyc ontology, for example, is organized around an extensive library of such *micro-theories* [17]. Modularity in ontologies may be pursued from the initial design or be imposed *post hoc*.

Although ontologies all broadly target the representation of domains of knowledge, the formal representation of specific ontologies often differ with respect to expressivity. Typically, top level ontologies will employ variants of first-order logic, while larger, but shallower semantic web ontologies will employ variants of description logic. There is then the usual trade-off between expressivity and computational complexity. 'Light' versions of top level ontologies may consequently be provided for practical applications.

2.2 Heterogeneous Ontologies

The increasing diversity and complexity of ontology applications have made questions of best design practice increasingly important. The most traditional, although still often practised, approach is to assume that a single top level ontology should be sufficient to cover all domains. However, debate on just which top level ontology would suffice for such a task has been ongoing for over 30 years in AI and the proposals forthcoming are still primarily limited in generality to particular domains or application areas.

An alternative approach is to accept that different communities and different areas of knowledge may require distinct modeling strategies. This latter approach, relying on what are termed *heterogeneous ontologies*, has several advantages. First, It constitutes a logical extension of the notion of modular ontologies – whereas traditional modular ontologies are composed of subtheories expressed in a common formal language and defined with respect to a common set of entities, heterogeneous ontologies allow modules to vary not only with respect to the domains that they are modeling but also with respect to the logical expressivity and logical languages adopted. Second, reasoning with, and maintenance of, complex ontologies is improved by virtue of the greater independence of components.

In previous work in the area of spatial representations, we developed techniques particularly supportive of heterogeneity [4]. Entire families of formal calculi exist for qualitative reasoning about space and specific reasoning tools have been developed that maximize performance for particular classes of spatial formalization. As a consequence, any enforcement of a common formalization within a single logic would not only have demanded a homogeneity unmotivated by the domains at hand but also have compromised the many fine-grained tailored reasoning solutions available. Building on this foundation, a general formal framework has was specified and implemented within which heterogeneous modules couched in distinct logics and drawing on diverse reasoning support can be combined into single coherent specifications.

Central to this framework is the the notion of a *hyperontology* [16], a concept defined and subsequently embodied within the *Distributed Ontology Language* (DOL) specification, accepted as an OMG standard in 2016 [21]. DOL specifications allow ontology modules to be specified in any logic for which a translation mechanism has been defined and then be freely combined for the purposes of hybrid reasoning. Permitting combinations of this kind generalizes beyond alignment approaches by following 'logic translations' that flexibly map between theories [15]. Multiple, overlapping and even conflicting formal specifications can thus be maintained simultaneously. A growing collection of heterogeneous ontologies, ontology components and inter-ontology mappings can be found in the OntoHub semantic repository [7]. Heterogeneous ontologies offer several important capabilities that we are now exploring in the context of defining ontologies for robotics, as we shall now see.

2.3 Heterogeneous Ontologies from a Robotics Perspective

As robots start to accomplish human-scale tasks in everyday environments, the need for knowledge representation and processing increases substantially due to both the complexity of such environments and the lack of well-structured definitions of tasks. Even when an ontology that covers all of the relevant concepts for a corresponding domain can be created and used, the cognitive architectures of robots need to map these concepts to other non-physical or cognitive components of the robot. This has several important consequences.

First, robot ontologies [25] should be enriched in terms of relations such as synonyms and similarity indices. This is mostly needed for interaction with humans and processing task definitions. Since communication using natural language involves considerable ambiguity and implicit information, ontologies for robotic applications need to be sufficiently comprehensive to resolve these issues correctly. Second, robots have to represent action concepts in terms of control programs and parameters. For instance, an action called *power grasp* should correspond to grasp controllers and, since it is a special type of grasping, parameters such as applied force must be bound to certain intervals. Similarly, perception concepts need to be captured in terms of perception programs and parameters. Third, personal robots need to have mechanisms for logging their experiences in order to use them to improve their skills. In order to make use of such experiences in reasoning and learning, they must be represented in terms of the same concepts that the robot ontologies have. And fourth, reasoning for robots often includes a planning, or envisioning, or simulation component, which is useful to either verify existing behaviors or create new ones. Similar to experiences, robotic simulations need to use the same concepts for initializing scenarios and generating datasets that the robot will use during actual execution for reasoning and learning purposes.

Figure 2 illustrates an example of how our robotic application already makes use of heterogeneity. In this example, the Cyc base ontology is extended with concepts from the intelligent robotics domain such as the *Robot* concept and *capableOf* relation [14]. As seen from the figure, *PR2* (right) is a robot capable of *pouring*. Some specific robot, *PR2_9su6*, may then first simulate a pouring action, *Pouring_x76a*, with a set of parameters that its planning algorithm deems suitable. After it sees that the simulated behavior has the expected effect, it actually executes a pouring action, *Pouring_4gp1*, with a similar parameter set. During this execution, the robot perceives *cup_x143* with its perception module and, using the semantic map of the environment and 6d pose of the cup, reasons that that cup is actually on *table_a9kl*. Thus, a procedural attachment of *on-Physically* relation between the cup and the table is performed dynamically.

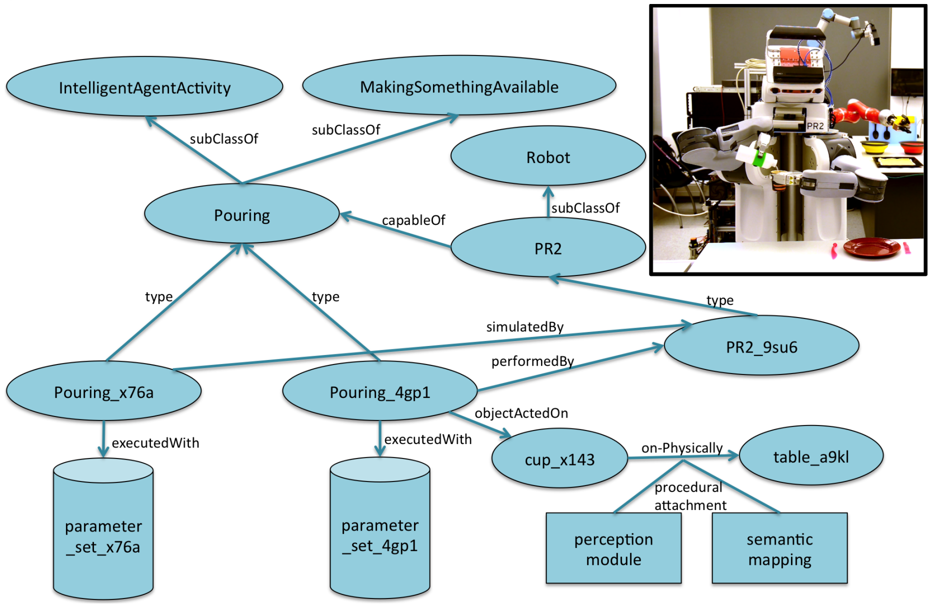

Fig. 2. A Robotic Application Extending Cyc Ontology Action Classes

Crucially, the various areas of concern – the 'pouring' micro-theories, simulation and perception – all have different properties but nevertheless need to be combined in the service of successful (and flexible) completion of the action. These kinds of requirements are then natural candidates for treatment in terms of heterogeneous ontologies. Whereas previous approaches to the use of ontology in robotics have focused on the traditional knowledge representation aspects of ontologies, native support for heterogeneity allows us to consider integrating other important functionalities for ontologies as well.

2.4 Heterogeneous Ontologies from a Language Perspective

Work on processing natural language has also raised several opportunities for beneficially employing heterogeneous ontologies. When analyzing or producing natural language, it is common to posit a level of relatively shallow semantics that is directly formed following principles of compositionality. This level is typically underspecified with respect to context in that it does not commit to specific world referents and may leave a variety of other aspects unresolved. Interpretation then requires an additional step of contextualization.

Consider, for example, the sentence and straightforward corresponding underspecified semantics given in example (1):

(1) a. Pour the water into the bowl.

b. $e : \mathrm{pour}'(e) \wedge \mathrm{actor}(e, h_1) \wedge \mathrm{actee}(e, w_1) \wedge \text{spatial-destination}(e, b_1)$

A semantics of this kind identifies the semantic type of the predicate at issue and the semantic role relationships between the predicate and further discourse entities introduced by the grammatical participants in the linguistic utterance.

Each of these further entities may similarly be given a semantic type. The task of contextualization then relates discourse entities to actual referents in a situation and fills in linguistically underspecified aspects of the description.

These distinct levels of representation are addressed in terms of heterogeneous ontologies in the work of Bateman, Hois, Ross and Tenbrink [5]. Here an extensive linguistically motivated ontology of semantic types and relations was defined for the shallow semantics produced by compositional analysis. This ontology, called the *Generalized Upper Model*, is specified in description logic. Particular areas within this ontology are then related to modular subtheories responsible for further reasoning. As an example, the sentence (1a) makes reference to a particular spatial relationship, that of an object moving along a path ('into'). The sentence itself makes no commitment to the kind of object involved, nor to the details of the path. All that is specified is the end point of the path. Nevertheless, on the basis of this abstract specification, it is still known that following successful execution of the action ('pouring') the object ('water') is in a location functionally controlled by the destination ('bowl'). This is termed a *two-level* spatial semantics [3].

Heterogeneous ontological specifications offer support for precisely these kinds of multi-leveled representation. The shallow semantic ontology, the subtheories characterizing particular kinds of motion and spatial relationships, as well as the contextualized semantics anchored in specific physical situations are all allowed their own independent existence, drawing on diverse kinds of formalizations as required. Maintaining such distinct levels of specification permits much of the extreme flexibility observed in real uses of spatial language to be covered, avoiding premature commitments to interpretations that only follow in specific situations. In the Cyc extract above, for example, the supertype *MakingSomethingAvailable* is such an overcommitment, only applying to particular cases of 'pouring' and not 'pouring in general'.

This establishes a striking parallel with the heterogeneity requirements considered in our robotic application above. Work in cognitive linguistics and linguistic construction grammar increasingly argues for accounts drawing on notions of force dynamics [24] and simulation [9] for linguistic semantics in general. The availability of simulation, experience logs and ontological information within a heterogeneous specification for hybrid specifications for robotics therefore offers a promising framework for integrating these requirements in a beneficially complementary fashion. Thus, instead of relying solely on logical inference as formerly pursued within linguistic semantics to work out, for example, that the water in example (1) above ends up in the bowl, we can augment this both with experiential knowledge about what happens when some agent attempts this and with direct simulation following physical laws as suggested in Figure 2.

3 Capturing and using heterogeneity: an example

To emphasize the need for heterogeneous ontologies in the context of mobile robotics and everyday manipulation activities, we elaborate further on the action of pouring because it is well understood, highly context dependent (and thus knowledge intensive), and still challenging for modern robotics platforms; the requirements involved here generalize well to a rich variety of similarly complex actions. Despite this often under appreciated complexity, humans can pour with ease even in unknown environments and using previously unknown objects. Enabling a robot to perform pouring actions at the scale of human performance requires deep knowledge about the activity, the objects involved, and how to handle them.

Assume a humanoid robot capable of picking up objects in a human-like fashion. The actual performance of the pour action depends on many contextual factors such as the physical properties of the substance that is poured (e.g., thickness), from where it is poured (e.g., pot or glass), and to which destination it is poured. Some pouring variants may even require pouring the substance in a specific pattern, e.g., to evenly spread sauce on a pizza dough before baking.

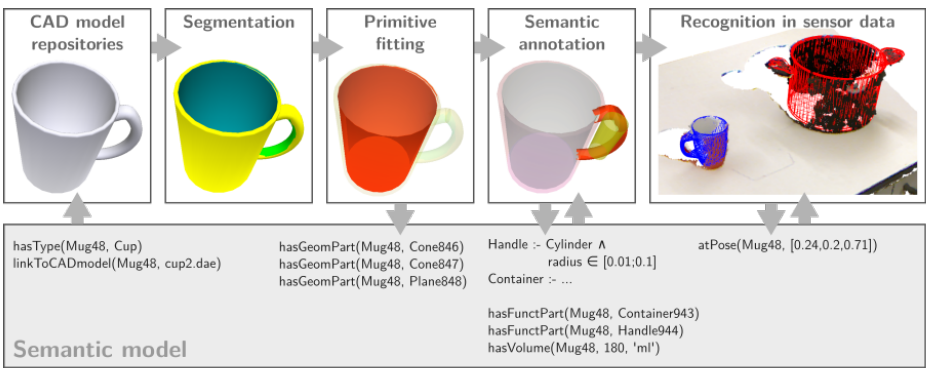

Fig. 3. Overview of the system for the semantic interpretation of geometric object models. Figure is adapted from Tenorth et al. [26].

One highly relevant example from the household domain is to pour dough mix onto a frying pan or a pancake maker in order to cook the dough. Let us assume the dough mix is inside of a container from which the robot can pour and that the action is performed within a perception-action loop in which the robot perceives the pose of the dough mix container and other visual properties such as its size. Let us further assume that the robot does not know the functional parts of the object in advance, and that the robot can find a mesh corresponding to the object in a mesh repository (by matching the visual features); a rich set of different mesh formats are widely supported, including *Collada* meshes which are represented using XML format. Information encoded in the mesh has no formal semantics and is not directly captured by common ontology formalisms, but nevertheless contains valuable information for task execution. Thus, we desire an (intermediate) representation of the information encoded in the mesh in terms of the robot's ontology. Using the algorithm described in [2], we can automatically find functional parts of the object, such as its handle, based on fitting primitive geometry with the mesh. This procedure is depicted in Figure 3. Thus, we can infer that $Container(?c) \wedge has\_part(?c, ?handle) \wedge Handle(?handle)$ by exploiting implicit knowledge encoded in the mesh.

A robot is often required to combine information from a variety of heterogeneous sources in order to draw the right conclusions. Light-weight simulation, where some actions such as navigation are abstracted away [19], can be used to verify the feasibility of some manipulation actions given a rough model of the environment (is a particular pose reachable? is a given location visible from another?). Also, the object's visual features are accessible to the robot through internal data structures encoded by the perception system. The format of such object designators is usually a nested list of key-value pairs where the keys correspond to feature names and the values to the feature values (without formal semantics). For generating an appropriate grasp point, for example, the robot has to combine knowledge about functional parts of the object with its perceived pose. Furthermore, spatial relations between involved objects are highly relevant during pouring because the container needs to be functionally *above-of* the pouring destination so that the poured substance is not spilled. Seen traditionally, the robot's ontology consequently needs to be augmented by a rule that implies this spatial relation based on some spatial calculus that internally accesses the data structures of the perception system. In terms of heterogeneity, this can equally be seen as a case of formally linking distinct areas of formalization.

Similarly, the robot may also need to reason about the effects and failures of actions and how to recover from them. Failure handling is usually explicitly encoded in plans without a formal causal model that describes, for example, that spilling is likely to happen if the source container is held too high. Such motion constraints can be inferred from experiential and simulated knowledge. Action effects concern individuals involved in the action and so cannot be expressed terminologically using OWL. In the case of pouring dough mix onto a pancake maker, the robot may want to reason about the cooking process that was started by the pouring action. Using rules we can express this relation between the action and the cooking process as follows:

```
HeatingProcess(?heating) ∧ causedBy(?heating, ?src) ∧
thermicallyConnectedTo(?src, ?food) ∧
EdibleStuff(?food) ∧ (not Cooked)(?food) →
   CookingFoodProcess(?cooking) ∧ causedBy(?cooking, ?src) ∧
   objectActedOn(?cooking, ?food) ∧
   processStarted(?heating, ?cooking)
```

This rule implies that a cooking food process starts when some edible stuff is thermically connected to a heating source which is the cause of a heating process, that the cooking process affects the edible stuff, and that the heating process is the cause for the cooking process. Involved objects are assumed to be thermically connected if one object is on top or inside of another. The existence of this predicate is computed on demand (cf. Figure 4) exploiting the heterogeneous data structures at hand (e.g., perception, action, sensors, semantic models). Formally linking the predicates given in such specifications with their simulated effects, perception and with previous experience logs opens up such semantics to a far more 'embodied' and 'experiential' level of representation and reasoning.

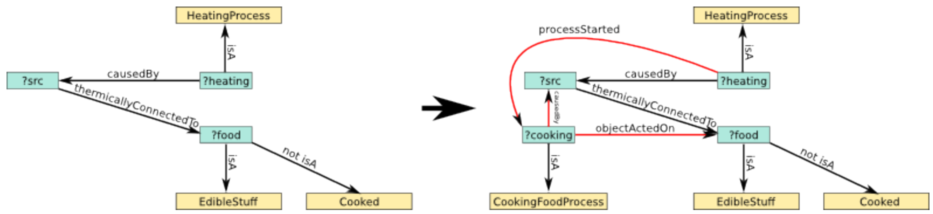

Fig. 4. A visual description of the cooking rule effects. Note that the rule requires new chains of properties to be established between particular pairs of individuals (?*food* and ?*heating*, ?*src* and ?*heating*) passing through a new individual (?*cooking*), which is not expressible in OWL restrictions.

4 Related Work

The need for ontologies and heterogeneous information sources is increasingly evident in robotics research. The Proteus project [18] employs ontologies to make explicit the semantics of models assumed by domain-specific languages, and to enable knowledge reuse. Krüger *et al.* propose *Object-Action complexes* (OACs) [13, 28], symbolic representations of sensorimotor experience and behaviors which formalize the connections between objects and associated actions. OACs have two main components: a symbolic description of an object (a prediction function defined over an attribute space, together with prediction reliability measures), and a specification for execution and learning of the OAC.

The affordance network ontology for robots (AfRob) [27] is another knowledge-enabled robotic framework that represents affordance relations of robots for an object dataset. The affordance concepts are also grounded in simple visual perception algorithms to identify certain shapes and spatial relations between objects/features. The affordances are used to assist object perception and scene interpretation, as well as to suggest grasps that are appropriate for the function of an object. AfRob also pushes beyond defining ontology concepts in purely logical terms, and is therefore an example of a heterogenous approach to ontology construction. In the work reported here we seek to improve the formalization of such heterogeneity.

Schlenoff *et al.* have initiated the IEEE-RAS working group *Ontologies for Robotics and Automation* (ORA) [23], whose goal is to develop standards for knowledge representation and reasoning in robotics. Although not a hybrid ontology, it is noteworthy for being by far the largest effort to create robotics ontologies to date, and the group has presented a proposal for a core ontology for robotics [22] (now adopted as an IEEE standard), and also produced ontologies for more specific use cases, such as industrial automation [10]. It uses well-established techniques for ontology development as well as previously established ontologies such as Pease's SUMO [20]. The benefit of more heterogeneity can also be seen in extensions such as the inclusion of position information, however: instead of simply adding these concepts below top-level concepts [10], the heterogeneous approach would allow more flexible access to distinct organizations of spatial information simultaneously, supporting multiple reasoning solutions as well as more modular development and maintenance.

5 Outlook

Standards and developments such as those mentioned in the previous section are necessary for the exchange of software and scientific results between institutes. They must also, however, by their very nature be relatively conservative in terms of the knowledge representation techniques employed. For the kinds of future research into flexible and intelligent robotic behaviors considered here, we believe it essential that more diverse sources of knowledge, experience and perception be related cleanly to the overall enterprise of pursuing ontological representations for robotics. An inherently heterogeneous framework of the kind proposed here is one step towards supporting such capabilities.

References

1. R. Arp, B. Smith, and A. D. Spear, *Building Ontologies with Basic Formal Ontology.* Cambridge, MA: MIT Press, 2015.
2. M. Attene, B. Falcidieno, and M. Spagnuolo, "Hierarchical mesh segmentation based on fitting primitives," *The Visual Computer*, vol. 22, no. 3, pp. 181–193, 2006.
3. J. A. Bateman, "Language and Space: a two-level semantic approach based on principles of ontological engineering," *International Journal of Speech Technology*, vol. 13, no. 1, pp. 29–48, 2010.
4. J. A. Bateman, S. Borgo, K. Lüttich, C. Masolo, and T. Mossakowski, "Ontological Modularity and Spatial Diversity," *Spatial Cognition and Computation*, vol. 7, no. 1, pp. 97–128, 2007.
5. J. A. Bateman, J. Hois, R. J. Ross, and T. Tenbrink, "A linguistic ontology of space for natural language processing," *Artificial Intelligence*, vol. 174, no. 14, pp. 1027–1071, September 2010.
6. S. Borgo and C. Masolo, "Ontological Foundations of DOLCE," in *Theory and Applications of Ontology: Computer Applications*, R. Poli, M. Healey, and A. Kameas, Eds. Dordrecht, Heidelberg, London and New York: Springer, 2010, pp. 279–296.
7. M. Codescu, E. Kuksa, O. Kutz, T. Mossakowski, and F. Neuhaus, "Ontohub: A semantic repository for heterogeneous ontologies," *CoRR*, vol. abs/1612.05028, 2016.
8. J. Euzenat, P. Shvaiko *et al.*, *Ontology matching.* Springer, 2007, vol. 18.
9. J. Feldman, "Embodied language, best-fit analysis, and formal compositionality," *Physics of Life Reviews*, vol. 7, pp. 385–410, 2010.
10. S. R. Fiorini, J. L. Carbonera, P. Gonçalves, V. A. Jorge, V. F. Rey, T. Haidegger, M. Abel, S. A. Redfield, S. Balakirsky, V. Ragavan, H. Li, C. Schlenoff, and E. Prestes, "Extensions to the core ontology for robotics and automation," *Robot. Comput.-Integr. Manuf.*, vol. 33, no. C, pp. 3–11, Jun. 2015.
11. T. R. Gruber, "A translation approach to portable ontology specifications," *Knowledge Acquisition*, vol. 5, no. 2, pp. 199–220, 1993.
12. N. Guarino, "The Ontological Level," in *Philosophy and the Cognitive Sciences*, R. Casati, B. Smith, and G. White, Eds. Vienna: Hölder-Pichler-Tempsky, 1994, pp. 443–456.
13. N. Krüger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Wörgötter, A. Ude, T. Asfour, D. Kraft, D. Omrčen *et al.*, "Object–Action Complexes: Grounded abstractions of sensory–motor processes," *Robotics and Autonomous Systems*, vol. 59, no. 10, pp. 740–757, 2011.

14. L. Kunze, T. Roehm, and M. Beetz, “Towards semantic robot description languages,” in *Robotics and Automation (ICRA), 2011 IEEE International Conference on.* IEEE, 2011, pp. 5589–5595.
15. O. Kutz, C. Lutz, F. Wolter, and M. Zakharyaschev, “$\mathcal{E}$-Connections of Abstract Description Systems,” *Artificial Intelligence*, vol. 156, no. 1, pp. 1–73, 2004.
16. O. Kutz, T. Mossakowski, and D. Lücke, “Carnap, Goguen, and the Hyperontologies: Logical Pluralism and Heterogeneous Structuring in Ontology Design,” *Logica Universalis*, vol. 4, no. 2, pp. 255–333, 2010, special Issue on ‘Is Logic Universal?’.
17. D. Lenat, “Cyc: a large-scale investment in knowledge infrastructure,” *Communications of the ACM*, vol. 38, no. 11, November 1995.
18. G. Lortal, S. Dhouib, and S. Gérard, *Integrating Ontological Domain Knowledge into a Robotic DSL.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 401–414.
19. L. Mösenlechner and M. Beetz, “Fast temporal projection using accurate physics-based geometric reasoning,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, May 6–10 2013, pp. 1821–1827.
20. I. Niles and A. Pease, “Toward a Standard Upper Ontology,” in *Formal Ontology in Information Systems (FOIS 2001)*, C. Welty and B. Smith, Eds. Amsterdam: IOS Press, 2001, pp. 2–9.
21. OMG, “The Distributed Ontology, Modeling, and Specification Language (DOL),” OMG, Tech. Rep., 2015.
22. E. Prestes, J. L. Carbonera, S. R. Fiorini, V. A. M. Jorge, M. Abel, R. Madhavan, A. Locoro, P. J. S. Gonalves, M. E. Barreto, M. K. Habib, A. Chibani, S. Grard, Y. Amirat, and C. Schlenoff, “Towards a core ontology for robotics and automation,” *Robotics and Autonomous Systems*, vol. 61, no. 11, pp. 1193–1204, 2013.
23. C. Schlenoff, E. Prestes, R. Madhavan, P. Goncalves, H. Li, S. Balakirsky, T. Kramer, and E. Miguelanez, “An IEEE standard ontology for robotics and automation,” in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on.* IEEE, 2012, pp. 1337–1342.
24. L. Talmy, “Force dynamics in language and cognition,” *Cognitive Science*, vol. 12, no. 1, pp. 49–100, January–March 1988.
25. M. Tenorth and M. Beetz, “KnowRob – A Knowledge Processing Infrastructure for Cognition-enabled Robots,” *Int. Journal of Robotics Research*, vol. 32, no. 5, pp. 566 – 590, April 2013.
26. M. Tenorth, S. Profanter, F. Balint-Benczedi, and M. Beetz, “Decomposing CAD Models of Objects of Daily Use and Reasoning about their Functional Parts,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo Big Sight, Japan, November 3–7 2013, pp. 5943–5949.
27. K. M. Varadarajan and M. Vincze, “Afrob: The affordance network ontology for robots,” in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on.* IEEE, 2012, pp. 1343–1350.
28. F. Wörgötter, A. Agostini, N. Krüger, N. Shylo, and B. Porr, “Cognitive agents – a procedural perspective relying on the predictability of Object-Action-Complexes (OACs),” *Robotics and Autonomous Systems*, vol. 57, no. 4, pp. 420–432, 2009.