This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Angleraud, Alexandre; Houbre, Quentin; Kyrki, Ville; Pieters, Roel

Human-Robot Interactive Learning Architecture using Ontologies and Symbol Manipulation

Published in:

Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2018

DOI:

[10.1109/ROMAN.2018.8525580](https://doi.org/10.1109/ROMAN.2018.8525580)

Published: 06/11/2018

Document Version

Peer reviewed version

Please cite the original version:

Angleraud, A., Houbre, Q., Kyrki, V., & Pieters, R. (2018). Human-Robot Interactive Learning Architecture using Ontologies and Symbol Manipulation. In *Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2018* (pp. 384-389). [8525580] IEEE.
<https://doi.org/10.1109/ROMAN.2018.8525580>

This is the accepted version of the original article published by IEEE.

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Human-Robot Interactive Learning Architecture using Ontologies and Symbol Manipulation

Alexandre Angleraud¹, Quentin Houbre¹, Ville Kyrki² and Roel Pieters¹

Abstract—Robotic systems developed for support can provide assistance in various ways. However, regardless of the service provided, the quality of user interaction is key to adoption by the general public. Simple communication difficulties, such as terminological differences, can make or break the acceptance of robots. In this work we take into account these difficulties in communication between a human and a robot. We propose a system that allows to handle unknown concepts through symbol manipulation based on natural language interactions. In addition, ontologies are used as a convenient way to store the knowledge and reason about it. To demonstrate the use of our system, two scenarios are described and tested with a Care-O-Bot 4. The experiments show that confusions and difficulties in communication can effectively be resolved through symbol manipulation.

I. INTRODUCTION

Modern-day assistive robots are actively being deployed in people's homes and care facilities in order to assess their capabilities in providing support [1]. Whether such support is the manipulation of objects (e.g., object pick-up [2]) or an interface between a person and her family (e.g., tele-presence robots [3]), the quality of user interaction is critical in the acceptance of robots into a person's everyday life. Very few people interact with robots on a day-to-day basis and the people that do, are typically trained to do so (e.g., factory workers, care personnel). The wide scale acceptance and benefit of robots for the everyday life and support can be ensured if the communication between human and robot is easy, intuitive and includes learning strategies. Such developments are already being made in the area of manufacturing. The traditional way of programming robots in industry is via proprietary software interfaces that require great experience and expertise from highly trained professionals. In recent years solutions to this problem are being proposed by offering interfaces that can easily program the task of a robot. Examples of such systems that can generate sequential tasks on a manufacturing line are the robots Sawyer³ and Universal Robots⁴. These solutions are a great development in structured environments and in structured tasks that can be well defined, i.e., in production and manufacturing environments. However, when the communication between the robot and the human is less clear, and can contain ambiguity, alternatives have to be sought.

¹Faculty of Engineering Sciences, Cognitive Robotics, Tampere University of Technology, Finland alexandre.angleraud@tut.fi, quentin.houbre@tut.fi, roel.pieters@tut.fi

²School of Electrical Engineering, Intelligent robotics, Aalto University, Finland ville.kyrki@aalto.fi

³<http://www.rethinkrobotics.com/sawyer/>

⁴<https://www.universal-robots.com/>

Fig. 1. Grounding of a new symbol (spaghetti) in the knowledge base of the Care-O-Bot 4 using human-robot interaction through natural language.

For example, a person that is not trained to interact with a robot might not be aware of the proper terminology, or might not know which format to use in communication. Additionally, verbal instructions suffer from the fact that many synonyms can be used for an identical concept (e.g., the use of the word 'mug' and 'cup' is interchangeable). Even specifying tasks indirectly implies certain locations that might not be known (e.g., 'get me a drink', implies the kitchen as location). Programming in advance all possible communication means and synonyms, and integrating fall back mechanisms when confusion arises is tedious and decreases the ease of use, and ease of access to robots.

In this work we propose a solution to this communication problem with a system that allows to handle unknown concepts by using the current knowledge of the robot. Information is delivered through a natural language processing module that provides a ($\langle \text{action} \rangle, \langle \text{target} \rangle$) pair for reasoning before storing the newly acquired data in ontologies. Our interlocutor for communication is a Care-O-Bot 4. We demonstrate the results with several typical scenarios in which queries from human to robot are given (see Figure 1).

II. RELATED WORK

Human-robot interaction for assistance and care has been an ongoing topic of research for several decades [4] and one of the main drivers of this development is the aging population. More precisely, the so called 'aging in place' [1], for which robotic assistance could provide an extended independence to elderly people. Besides this support in one's home or in assisted living environments, dedicated

interaction with robots is also being developed for manufacturing environments [5]. What these two seemingly different human-robot interaction scenarios have in common is their respective goals: supporting people in their daily life, i.e., for care or for work.

Foremost, interaction should be safe and intuitive. This implies that a robot should be aware of its environment, its capabilities and the risks that accompany these. Such awareness is inherent in humans, but typically has to be programmed for robots. Much research exists in providing robots with capabilities for situational awareness, with typical examples such as navigation [6] and semantic mapping [7]. The capability that can lift the awareness to a higher level is reasoning. Reasoning over the knowledge that a robot contains allows for inference, deduction and conclusion, which would not be possible otherwise.

In [8] a cognitive architecture as well as a knowledge model are described that aim at offering robots artificial cognitive capabilities. Developed for social human-robot interaction this extensive piece of work includes, among others, geometric reasoning and situation assessment, knowledge acquisition and representation and multi-modal dialogue. Reasoning relies on the OpenRobots Ontology (ORO) system [9], which represents knowledge in the first-order logic formalism as triples following the Resource Description Framework (RDF). In this model a triple is typically in the form of subject-predicate-object.

Following the idea to reason over RDF components, Tenorth and Beetz describe in [10] the KnowRob system that is specifically designed to provide autonomous robots with the knowledge needed for performing everyday manipulation tasks. KnowRob is based on Prolog, which internally stores knowledge in terms of Description Logic and provides the vocabulary for describing knowledge about actions, events, objects, spatial and temporal information. Extensive demonstrations are shown for complete domestic tasks such as cooking pancakes and serving drinks. Based on KnowRob, projects such as RoboEarth [11] focus on making robots capable of learning from shared experiences.

In this paper however, learning is studied through human-robot interactions. More precisely, we study the possibility to teach new or related concepts to a robot by communicating and manipulating symbols in a similar way as the Object Action Complex (AOC) in [12]. To our knowledge very few works are tackling it from this angle. In [13] a cognitive architecture is developed to allow the learning of objects and actions. While in [14] the teaching process also uses semantic information, the design of the teaching process is different and we tried to adopt a communication model as close as possible from what the end user could expect. This is why the core of the conversation uses a model inspired by human-to-human conversations similar to [15]. As for the reasoning, we extend the KnowRob system on two aspects. First, by creating a custom set of ontologies using the Protégé editor [16] and integrating them in the knowledge base. Second, by developing new Prolog queries interacting with our ontologies to deal with symbol grounding [17].

III. ARCHITECTURE

The conceptual reasoning system is represented by three layers, i.e., the input layer, the reasoning layer and the action layer (see Figure 2). The input layer represents the communication between the environment and the robot and is guided by a state machine (see Figure 3). The reasoning layer forms the core of the system by providing various reasoning services that interact with the knowledge base. The knowledge base itself contains information on the environment, the robot and its capabilities and local objects. Reasoning extends and updates the knowledge base with actions and targets that are taught by a person. Finally, the action layer serves as output to the conceptual reasoning. In this chapter we will explain in more details what happens in the input and reasoning layer. The following entities will be considered.

Requests are the set of commands that a user is able to send. A Request R is defined as :

$$R = (A, T), \quad (1)$$

where A refers to the symbol of an action and T refers to the symbol of a target.

Actions are the potential commands that the robot can send to the action layer to perform a motion, or speak out a question for instance. An action A takes an optional target T as parameter and can be composed of a number N of sub actions. As a result :

$$A = \{a_i(t_i)\}_{i=1}^N \quad (2)$$

where a_i and t_i are the sub action and sub target of the step i for the action A .

Targets are notions of concrete elements from the environment. To a target T is associated a number N of properties, thus :

$$T = \{p_i\}_{i=1}^N \quad (3)$$

where p is a property of the target attached to a primitive data type value (a string or an integer for instance).

A. Input layer

The Input layer is in charge of delivering meaningful tokens to the reasoning module from natural language inputs provided by the user. Considering the complexity of natural language processing, we adopted a pattern matching approach allowing us to classify each utterance in a finite number of categories. More precisely, once the inputs have been converted to text and the different words turned into tokens, we use a processing step called chunking⁵. In this step, instead of defining the utterances to be recognized by the robot we define what should be ignored. We then apply a set of rules on the remaining tokens that will define the next state of the robot. To each state is attached a different set of rules which guides the shape of the input expected from the user. Figure 3 shows the state machine attached to the speech recognition module. The calls from the input layer

⁵<http://www.nltk.org/book/ch07.html>

Fig. 2. Overview of the system. The input layer detects spoken requests and transforms them into a comprehensive set of tokens for the reasoning layer. There, logic reasoning is performed through Prolog and interaction with the ontologies constituting the knowledge base. If possible the spoken request is converted into an action command sent to the robot and the specificities of the action are dealt with in the action layer. ROS is used to build the communication architecture, here represented with the arrows.

to the functionalities of the reasoning layer are contained in the different states.

The default state is the `Listening` state. It is the initial state as well as the returning state after an action has been performed or a word has been taught. Whenever a request is fully understood the system goes into a `Action_Execution` state. It can be seen as a blocking state for the speech recognition, allowing the robot to finish its current action before another request can be taken into consideration.

The two other states, `Grounding_Target` and `Grounding_Action` are triggered by the addition of the word "teach" before the object of the teaching. Moreover, the word "action" between "teach" and the object of the teaching will trigger `Grounding_Action` whereas without it, the system evolves towards `Grounding_Target`. For instance "I will teach you the action bring a cup" will trigger the state `Grounding_Action`, when "I will teach you what mug is" will trigger `Grounding_Target`. They both differ from "Bring me my mug" that would be interpreted as a request and not as the starting point of the teaching process.

For teaching the actions, a sequence of requests composed of the same (`<action>`, `<target>`) pair as the initial `Listening` state is expected. They have to be already known by the robot and represent the components of the action that is being taught. As for the targets, only a sequence of reference concepts is expected. It is sufficient to use a word describing the mother class of the new target. However the user can also use an equivalent concept to describe it using the keyword "like" before the reference word (e.g., "A mug is like a cup").

When the utterance does not respect the previous schema, the straightforward approach is to evolve to an `Error` state forcing the robot to abandon the learning. Nevertheless for having easier interactions we loop over the `Grounding`

Fig. 3. State machine of the Speech Recognition module. Each state has its own rules that guide the type of sentences that the robot expects to hear.

state. Finally, to provide a clear feedback, each transition from a state to another is associated with an oral statement from the robot. For instance "I am listening" after "I will teach you what mug is"

B. Reasoning layer

The second layer is where the reasoning takes place. It receives as input the tokens delivered by the natural language processing module that we described previously and forwards an action command or sends back an unknown token message. The different components of this layer are as follows (Figure 4 illustrates this layer).

Knowledge base. It is composed of information describing the environment, the robot, the different actions and targets. This information is stored in ontologies. It is technically a separate entity from the robot and the reasoning layer as the knowledge base can be accessed from any digital platform due to the semantic web technology.

Prolog libraries. They are the link between the data and the rest of the program and hence are called before each

Fig. 4. Components of the reasoning layer. Its goal is to analyze the tokens of a spoken request to determine if it is able to perform it or not. It uses a Prolog based reasoning and the data model to do so. It can also detect if more information is needed to understand the request. In which case the corresponding grounding server is started. Finally, it can output an action command for the action layer. The arrows represented here symbolize ROS communication.

read or write operation in the ontologies. As a rule based language, Prolog allows us to perform logic reasoning over the information stored in the knowledge base.

Grounding modules. They are the two entities that get activated when a teaching state is reached in the input layer.

Kernel. It is the main entity, its role is to distribute the tasks to the other components and sends the output of the reasoning layer. This is where the logic of the interaction in the semantic web described in [15] is implemented by performing a syntactic analysis followed by a semantic analysis on each entry received from the input layer.

Data model. It mirrors how the information is stored in the knowledge base. In other words it acts as a template that is used to shape the information from the code to the knowledge base or process it from the knowledge base to the code.

IV. DEMONSTRATIONS

The aim of the scenarios is to use human-robot interaction for conceptual reasoning. Therefore, relating existing concepts and learning new concepts are human-guided by letting the person know when a concept is unknown. After the notification has been sent the user can decide to teach the word or to send another request formulated in another way. If a teaching process is started the aim is to interactively resolve the communication ambiguity using the patterns described previously. To evaluate this, two scenarios are devised in an assistive setting in which a person interacts with the service robot. The experiments have been carried out in a laboratory setting and the person interacting with the robot understood the system. Figures 5 and 6 show different time frames during the interactions and the following section describes the two scenarios and their developed outcomes. Being an

open source project, our developments are available for the robotics community.⁶

The robot used for human-robot interaction experiments is the Care-O-Bot 4⁷. It provided the two actions used during our scenarios, i.e. speech synthesis and movement of the base. To continue with the policy of the robot software, the systems are build on ROS and uses python together with Prolog as main languages. As for the speech to text functionalities, they use the Google recognition engine. Figure 2 shows the global organization of the system.

A. Cook pasta

The first HRI experiment describes how an unknown concept can be taught to a robot. We consider a human requesting from the robot to 'Cook spaghetti'. First, the robot does not know the action symbol ("Cook") and a conversation solving this ambiguity takes place. Next, a second conversation defining "spaghetti" as a sub concept of pasta is also needed. Finally, grounding equivalent concepts is also shown by asking to cook macaroni as the process to cook macaroni or spaghetti is similar.

TABLE I
SCRIPT OF THE DEMONSTRATION 1. IT INTRODUCES 3 DIFFERENT TYPES OF GROUNDING : ACTION, HIGHER LEVEL CONCEPT TARGET AND EQUIVALENT TARGET.

Utterance	Resulting state
"Cook spaghetti"	Listening
"I will teach you the action to cook pasta"	Grounding_Action
"Boiling pasta"	Grounding_Action
"Done"	Listening
"Cook spaghetti"	Listening
"I will teach you what spaghetti are"	Grounding.Target
"Spaghetti are pasta"	Grounding.Target
"Done"	Listening
"Cook macaroni"	Listening
"I will teach you what macaroni are"	Grounding.Target
"macaroni are like spaghetti"	Grounding.Target
"Done"	Listening

Table I describes the set of utterances submitted to the robot and the effects that they had on the state machine of the speech recognition module. For simplicity, it does not show what the robot answers. Note that the request "cook spaghetti" results in the same `Listening` state several times, first because the robot asked to define "cook" and then asked about "spaghetti". This first example shows the three different groundings possible in the system, i.e., an action, and for the targets a higher-level concept as well as an equivalent concept. We chose to define cook through only one action and related spaghetti using only one word, however, it would be possible to use more than those. The newly acquired knowledge is immediately reusable, as when the robot answers that it cannot cook spaghetti it means that the request is now understood, and only the physical capabilities are missing. Practically, acquiring knowledge in our system is a two steps process. First, a new class is

⁶<http://human-robot-interactive-learning.readthedocs.io>

⁷<http://www.mojin-robotics.de/>

Fig. 5. In scenario 1 the robot is asked to cook pasta, an action that is unknown to him. It thus asks about the ambiguous symbol. Using a sequence of sub actions already known, i.e., boiling water in this case, we can teach the meaning although the physical action would not be possible to perform.

Fig. 6. In scenario 2 the robot is asked to move to the sofa, a target that is unknown to him. It thus asks about the ambiguous symbol. Using a higher level concept already known, i.e., the room containing the sofa in this case we can solve the ambiguity of the request.

generated with the corresponding relations to the existing items in the knowledge base. Here for a example *Cook* is added as a subclass of *Action* and linked to *Boiling* with a relation *hasStep*. As for *Macaroni*, it is added as an equivalent class of *Spaghetti*. The second steps consists in generating individuals or in other words instances, corresponding to the new classes for the needs of the task execution.

B. Go to the sofa

The second experiment shows how an unknown concept can be converted into an action command after the ambiguity around its meaning has been resolved. We consider a human requesting from the robot to "go to the sofa", the robot is again asking to define the concepts. Table II describes the set of utterances submitted to the robot and the effects that they had on the state machine of the speech recognition module.

With respect to this scenario, "Move" is an action that was available on the Care-O-Bot and the position of "Kitchen" that we used to locate "Sofa" had been preloaded in the

knowledge base. Hence, *Go* is added as a subclass of *Action* and linked to *Move* with a relation *hasStep*. As for *Sofa*, it is added as an equivalent class of *Kitchen*. Therefore, once the two symbols were explained and individuals (or instances) created, the robot was able to enter the *Action_Execution* state and move to the location stored in the knowledge base. Table II describes the utterances and their effect with respect to this scenario.

V. DISCUSSION

To start with, two big assumptions were made about the interactions. First, that the user is giving the correct answer all the time, and second, that he/she always gives an answer. Additionally, the quality of the input depends largely on the quality of the microphone used.

The speech recognition module functioned satisfactory when keeping inputs relatively simple, i.e., short and concise sentences but even then the time between the spoken request and the end of the processing is still relatively long. Nonetheless, it uses a natural way of communication for humans and thus assumes no trained or qualified personnel.

The *Error* state, although not used per se in this first iteration is an important part of the system as it could break the interaction and with it the acceptance of the robot if not handled properly.

A concept can be an action or a target, and reasoning over the concepts is done separately. It is questionable if learning unknown concepts is truly learning or if it only consists of generating a sequence of already known concepts or (*<action>*, *<target>*)-pairs. In our view, learning means acquiring knowledge that was previously not known, whether or not underlying mechanisms or low level tasks are known. New concepts are introduced and new relationships between concepts are made, both within the ontology. This allows

TABLE II

SCRIPT OF THE DEMONSTRATION 2. IT SHOWS THE POTENTIAL TRANSITION TO THE ACTION LAYER AFTER THE AMBIGUITY AROUND THE LOCATION OF SOFA IS SOLVED.

Utterance	Resulting state
"Go to the sofa"	Listening
"I will teach you the action to go"	Grounding_Action
"Move to the place"	Grounding_Action
"Done"	Listening
"Go to the sofa"	Listening
"I will teach you where sofa is"	Grounding_Target
"Sofa is in the kitchen"	Grounding_Target
"Done"	Listening
"Go to the sofa"	Action_Execution

for new knowledge to be acquired by reasoning over the ontological concepts. The main advantage of this type of reasoning is that it brings predictability to the system and predictability is a key factor for the acceptance of robots to the general public [18].

Robotic applications have very specific demands regarding the abstract concepts in the knowledge bases that are hard to meet. One of the main challenges, the so call grounding problem [17] is to link an abstract knowledge representation and a particular control system of a robot. The semantic web technology accessed here through KnowRob offers the possibility to develop common-sense knowledge and reason about human or robot activities. In addition our system offers a bridge between symbols transmitted orally and physical actions to be performed by the robot.

Moreover, the architecture developed in this paper opens the door for additional features such as, extension of the state machine that could extend the range of actions that the robot is able to perform. To achieve this, additional work is necessary for the natural language processing module to bring more flexibility to our current model.

Finally, the system of tokens is convenient due to its extendability. For instance, if a single agent - single user scenario has been our first focus, adding other perception inputs could allow for multi users scenarios in a close future by assigning specific information to specific users for instance. In this paper we also focused on the high level decision making abilities of the robot but it is only a first step before including lower level mechanisms to make the transition between the reasoning layer and the action layer.

VI. CONCLUSIONS

The technical contribution of this work is the interactive, conceptual reasoning over actions and targets in an ontology. Communication between humans is often prone to misunderstanding and ambiguity due to differences in terminology, culture and age. Similarly, these difficulties extend to human-robot communication as well. With this work we aim to offer methods to communicate, resolve ambiguities and learn new knowledge in an interactive way. Instead of programming all possible interaction in advance, the robot should learn through interaction and be capable of reusing this knowledge in the future. The developed architecture consists of three layers. The input layer processes spoken requests and transmits an <action>,<target>-pair to the reasoning layer. There, the request is analyzed through logic reasoning to determine if it is understood or not and if the robot is able to perform what is being asked. The action layer deals with the specificities of the task. As a result the robot is treated more as an interactive partner than as a support tool. We believe that this is the future role of service robots. Two human-robot interaction scenarios with a Care-O-Bot 4 demonstrate our approach. Future work will focus on improving and extending our approach with respect to communication: instead of a robot confirming a learned result by replying 'Done', the robot should confirm the learned phrase itself. Additionally, we will assess our developments

with experiments outside the laboratory. On the one hand this will include HRI scenarios with people unfamiliar with robots, and on the other hand extend our approach to the domain of industrial human-robot collaboration.

ACKNOWLEDGMENT

The work presented here is financed by the Academy of Finland project: "Competitive funding to strengthen university research profiles", decision number 310325.

REFERENCES

- [1] D. Fischinger, P. Einramhof, K. Papoutsakis, W. Wohlkinger, P. Mayer, P. Panek, S. Hofmann, T. Koertner, A. Weiss, A. Argyros, and M. Vincze, "Hobbit, a care robot supporting independent living at home: First prototype and lessons learned," *Robotics and Autonomous Systems*, vol. 75, pp. 60–78, Jan. 2016.
- [2] *Personal Assistants: Emerging Computational Technologies*, ser. Intelligent Systems Reference Library, Cham, 2018, vol. 132.
- [3] A. Kristofferson, S. Coradeschi, and A. Loutfi, "A Review of Mobile Robotic Telepresence," *Advances in Human-Computer Interaction*, vol. 2013, pp. 1–17, 2013.
- [4] I. Leite, C. Martinho, and A. Paiva, "Social Robots for Long-Term Interaction: A Survey," *International Journal of Social Robotics*, vol. 5, no. 2, pp. 291–308, Apr. 2013.
- [5] G. Michalos, S. Makris, J. Spiliotopoulos, I. Misios, P. Tsarouchi, and G. Chrysolouris, "ROBO-PARTNER: Seamless Human-Robot Cooperation for Intelligent, Flexible and Safe Operations in the Assembly Factories of the Future," *Procedia CIRP*, vol. 23, pp. 71–76, 2014.
- [6] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, Dec. 2013.
- [7] R. B. Rusu, "Semantic 3d Object Maps for Everyday Manipulation in Human Living Environments," *KI - Künstliche Intelligenz*, vol. 24, no. 4, pp. 345–348, Nov. 2010.
- [8] S. Lemaignan, M. Warnier, E. A. Sisbot, A. Clodic, and R. Alami, "Artificial cognition for social humanrobot interaction: An implementation," *Artificial Intelligence*, vol. 247, pp. 45–69, 2017.
- [9] S. Lemaignan, R. Ros, L. Mösenlechner, R. Alami, and M. Beetz, "ORO, a knowledge management platform for cognitive architectures in robotics," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010, pp. 3548–3553.
- [10] M. Tenorth and M. Beetz, "KnowRob: A knowledge processing infrastructure for cognition-enabled robots," *The International Journal of Robotics Research*, vol. 32, no. 5, pp. 566–590, Apr. 2013.
- [11] M. Waibel, M. Beetz, J. Civera, R. D'Andrea, J. Elfving, D. Gálvez-López, K. Häussermann, R. Janssen, J. Montiel, A. Perzylo, B. Schieße, M. Tenorth, O. Zweigle, and R. De Molengraft, "RoboEarth," *IEEE Robotics & Automation Magazine*, vol. 18, no. 2, pp. 69–82, June 2011.
- [12] C. Geib, K. Mourao, R. Petrick, N. Pugeault, M. Steedman, N. Krueger, and F. Worgotter, "Object Action Complexes as an Interface for Planning and Robot Control," in *IEEE-RAS Humanoids Workshop: Towards Cognitive Humanoid Robots*, 2006, p. 7.
- [13] M. Scheutz, E. Krause, B. Oosterveld, T. Frasca, and R. Platt, "Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 1378–1386.
- [14] V. Paléologue, J. Martin, A. K. Pandey, A. Coninx, and M. Chetouani, "Semantic-based interaction for teaching robot behavior compositions," in *Robot and Human Interactive Communication (RO-MAN)*, 2017 26th IEEE International Symposium on, 2017, pp. 50–55.
- [15] P. Heim, T. Schlegel, and T. Ertl, "A model for human-computer interaction in the semantic web," in *Proceedings of the 7th International Conference on Semantic Systems*, 2011, pp. 150–158.
- [16] M. A. Musen, "The protégé project: a look back and a look forward," *AI Matters*, vol. 1, no. 4, pp. 4–12, June 2015.
- [17] S. Harnad, "The Symbol Grounding Problem," *Physica D*, vol. 42, pp. 335–346, 1990.
- [18] K. Dautenhahn, S. Woods, C. Kaouri, M. Walters, Kheng Lee Koay, and I. Werry, "What is a robot companion - friend, assistant or butler?" 2005, pp. 1192–1197.