4th International Symposium on New Frontiers in Human-Robot Interaction
21–22 April 2015, Canterbury, UK
M. Salem, A. Weiss, P. Baxter, K. Dautenhahn (Eds.)

Robot Learning from Verbal Interaction: A Brief Survey

Heriberto Cuayáhuitl[1]

Abstract. This survey paper highlights some advances and challenges in robots that learn to carry out tasks from verbal interaction with humans, possibly combined with physical manipulation of their environment. We first describe what robots have learnt from verbal interaction, and how do they do it. We then enumerate a list of research limitations to motivate future work in this challenging and exciting multidisciplinary area. This brief survey points out the need of bringing robots out of the lab, into uncontrolled conditions, in order to investigate their usability and acceptance by end users.

1 INTRODUCTION

Intelligent conversational robots are an exciting and important area of research because of their potential to provide a natural language interface between robots and their end users. A learning conversational robot can be defined as an entity which improves its performance over time through verbally interacting with humans and/or other machines in order to carry out abstract or physical tasks in its (real or virtual) world. The vision of such kinds of robots is becoming more realistic with technological advances in artificial intelligence and robotics. The increasing development of robot skills presents boundless opportunities for them to perform useful tasks for and with humans. Such development is well suited to robots with a physical body because they can exploit their input and output modalities to deal with the complexity of public spatial environments such as homes, shops, airports, hospitals, etc. A robot learning from interaction, rather than a robot that does not learn, is particularly relevant because it is not feasible to pre-program robots for all possible environments, users and tasks. Even though many robotic systems can be scripted or programmed to behave just as expected, the rich nature of interaction with the physical world, or with humans, demands flexible, adaptive solutions to deal with dynamic, previously unknown, or highly stochastic domains. Therefore, robots should be able to refine their already learned skills over time and/or acquire new skills by (verbally) interacting with its users and its spatial environment. An emerging multidisciplinary community at the intersection of machine learning, human-robot interaction, natural language processing, robot perception, robot manipulation and robot gesture generation, among others, seeks to address challenges in realising such robots capable of interactive learning.

This paper will provide a brief survey on robots that learn to acquire or refine their verbal skills from example interactions using machine learning. Conversational robots that draw on hand-coded behaviours, or robots learning from non-verbal interaction [3, 14], are therefore considered out of scope here.

[1] Heriot-Watt University, United Kingdom, email: hc213@hw.ac.uk

2 ADVANCES

2.1 What have robots learnt from conversational interaction?

The following list of representative conversational robots shows a growing interest in this multidisciplinary field, see Figure 1.

- The mobile robot *Florence* is a nursing home assistant [20, 17]. The tasks of this robot include providing the time, providing information about the patient's medication schedule and TV channels, and motion commands such as go to the kitchen/bedroom. The learning task consists in inducing a dialogue strategy under uncertainty, where the actions correspond to physical actions (motion commands) and clarification or confirmation actions. The robot's goal is to choose as many correct actions as possible.
- Iwahashi's non-mobile robot with integrated arm+hand+head learns to communicate from scratch by physically manipulating objects on a table [11]. The tasks of this robot include (a) acquisition of words, concepts and grammars for objects and motions; (b) acquisition of the relationships between objects; and (c) the ability to answer questions based on beliefs. The robot's goal is to understand utterances and to generate reasonable responses from a relatively small number of interactions.
- The mobile robot *SmartWheeler* is a semi-autonomous wheelchair for assisting people with severe mobility impairments [19]. The task of the robot is to assist patients in their daily locomotion. The learning task is similar as in the Florence robot, the induction of a dialogue manager under uncertainty, but with a larger state space (situations). The robot's goal is to reduce the physical and cognitive load required for its operation.
- A mobile robotic forklift is a prototype for moving heavy objects from one location to another [25]. Example commands include going to locations, motion commands, and picking up and putting down objects. The learning task consists in understanding natural language commands in the navigation and object manipulation domain. The robot's goal is to ground natural language commands (mapping commands to events, objects and places in the world [18]) in order to output a plan of action.
- The humanoid robot *Simon* manipulates physical objects on a table from human teachers [2]. The task of the robot includes pouring cereal into bowls, adding salt to salads, and pouring drinks into cups. The learning task is to ask questions to human demonstrators from three different types: label queries (Can I do it like this?), demonstration queries (Can you show me how to do it?) , and feature queries (Should I keep this orientation?). The robot's goal is to ask as good questions as possible in order to achieve fast learning from physical demonstrations.
- A KUKA mobile platform with manipulator ensembles simple furniture [24]. The task of the robot is to assemble IKEA furniture such as tables based on STRIPS-like commands. The learning

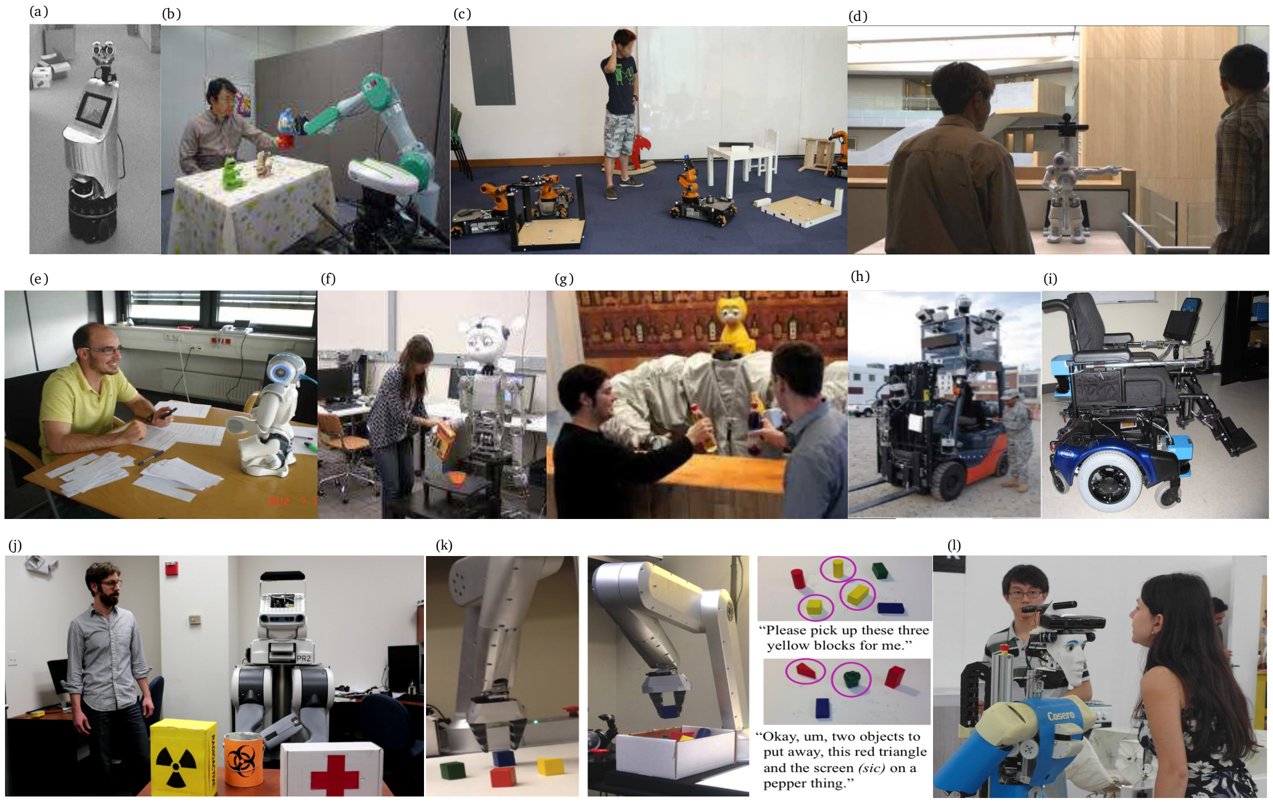

Figure 1. Example learning conversational robots: (a) Florence nursebot [20], (b) Iwahashi's robot [11], (c) Kuka furniture assembler [24], (d) Nao giving directions [1], (e) Nao playing quizzes [7], (f) Simon robot learning from demonstrations [2], (g) James bartender robot [12], (h) Forklift robot [25], (i) SmartWheeler [19], (j) PR2 learning new words [15], (k) Gambit picking up objects [16], and (l) Cosero receiving verbal commands [21]. See text in Section 2.

tasks consists in learning to ground language and to train a natural language generator in order to ask for help to humans (by generating words from symbolic requests) when the robot encounters a failure situation. The robot's goal is to ensemble furniture as independently as possible and to ask for help when failures occurred.

- The torso robot *James* serves drinks to people in a pub [12]. The task of the robot is to approach customers in natural language, to ask for the drinks they want, and to serve the requested drinks. The learning task consists in inducing a dialogue manager for multiparty interaction. The robot's goal is to serve as correct drinks as possible based on socially acceptable behaviour due to the presence of multiple customers at once in the robot's view.
- The humanoid robot *NAO* has been used to play interactive quiz games [7, 6]. The robot's tasks include engaging into interactions, asking and answering questions from different fields, and showing affective gestures aligned with verbal actions. The learning task consists in inducing a dialogue strategy optimising confirmations and flexible behaviour, where users are allowed to navigate flexibly across subdialogues rather than using a rigid dialogue flow. The robot's goal is to answer correctly as much as possible and to ask as many questions as possible from a database of questions.
- The humanoid robot *NAO* has been used to give indoor route instructions [1]. The task of the robot is to provide directions, verbally and with gestures, to places within a building such as offices, conference rooms, kitchen, cafeteria, bathroom, etc., based on a predefined map. The learning task is to induce a model of engagement to determine when to engage, maintain or disengage an interaction with the person(s) in front of the robot. The robot's goal is to direct people to the locations they are looking for.
- The mobile robot *PR2* has been used to acquire new knowledge of objects and their properties [15]. The tasks of the robot include to spot unknown objects, to ask how unknown objects look like, and to confirm newly acquired knowledge. The learning task is to extend its knowledge base of objects via descriptions of their physical appearance provided by human teachers. The robot's goal is to answer questions of its partially known environment.
- The robot arm *Gambit* has been used to study how users users refer to groups of objects with speech and gestures. The tasks of the robot is to move indicated objects in a workspace, via verbal descriptions of object properties and possibly including gestures. The learning task is to understand user intentions without requiring specialized user training. The robot's goal is to select, as correctly as possible, the referred objects on the table.
- The mobile robot *Cosero* has been used in the RoboCup at home competition, which has won several of them in recent years [21]. The tasks of the robot include to safely follow a person, to detect an emergency from a person calling for help, to get to know and recognise people and serve them drinks, and to bring objects from one location to another. The learning task is to extend its knowledge of locations, objects and people. The robot's goal is to carry out tasks autonomously—provided in spoken language—as expected and in a reasonable amount of time.

ID	Dimension / Reference	[20]	[11]	[19]	[25]	[2]	[24]	[12]	[7]	[1]	[15]	[16]	[21]	ALL
01	Learning To Interpret Commands	1	1	1	1	1	1	1	1	1	1	1	1	12
02	Dialogue Policy Learning	1	0	1	0	1	0	1	1	0	0	0	0	5
03	Learning To Generate Commands	0	1	0	0	0	1	0	0	0	0	0	0	2
04	Learning To Engage	0	0	0	0	0	0	1	1	1	0	0	0	3
05	Grammar Learning	0	1	0	0	0	0	0	0	0	0	0	0	1
06	Flexible Interaction	0	0	0	0	0	0	0	1	0	0	1	1	3
07	Speech-Based Perception	1	1	1	0	1	0	1	1	1	1	1	1	10
08	Language Grounding	0	1	0	0	0	1	0	0	0	0	1	0	3
09	Speech Production	1	1	1	0	1	0	1	1	1	1	0	1	9
10	Multimodal Fussion	0	1	1	0	1	0	1	0	1	1	1	1	8
11	Multimodal Fission	0	1	1	0	0	0	1	1	1	0	0	1	6
12	Multiparty Interaction	0	0	0	0	0	0	1	0	1	0	0	0	2
13	Route Instruction Giving	0	0	0	0	0	0	0	0	1	0	0	0	1
14	Navigation Commands	1	0	1	1	0	1	0	0	0	0	0	1	5
15	Object Recognition and Tracking	0	1	0	1	1	1	1	0	0	1	1	1	8
16	Human Activity Recognition	0	0	0	0	0	0	0	0	1	0	0	1	2
17	Localisation and Mapping	1	0	1	1	0	0	0	0	0	0	0	1	4
18	Gesture Generation	0	0	0	0	1	0	0	1	1	1	0	1	5
19	Object Manipulation	0	1	0	1	1	1	1	0	0	0	1	1	7
20	Supervised Learning	0	1	0	1	0	1	1	1	1	0	1	0	7
21	Unsupervised Learning	0	0	1	0	0	0	0	0	0	0	1	0	2
22	Reinforcement Learning	1	0	1	0	0	0	1	1	0	0	0	0	4
23	Active Learning	0	0	0	0	1	0	0	0	0	0	0	0	1
24	Learning From Demonstration	0	0	0	0	1	0	0	0	0	1	0	1	3
25	Evaluation w/Recruited Participants	1	0	0	0	1	1	1	1	1	0	1	1	8
26	Evaluation in Noisy/Crowded Spaces	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 1. Features of robots acquiring/using their verbal skills. While boolean values are rough indicators, real values are better indicators but harder to obtain.

2.2 How do conversational robots learn to interact?

Machine learning frameworks are typically used to equip robots with learning skills, and they differ in the way they treat data and the way they process feedback [13, 8]. Some machine learning frameworks addressed by previous related works are briefly described as follows:

- *Supervised learning* can be used whenever it comes to the task of classifying and predicting data, where the data consists of labelled instances (pairs of features and class labels). The task here is to induce a function that maps the unlabelled instances to labels. This function is known as a classifier when the labels are discrete and as a regressor when the labels are continuous. Conversational robots make use of classifiers to predict spatial description clauses [25], grounded language [11, 24], social states [12], dialogue acts [7], gestures [16], and engagement actions [1], among others.
- *Reinforcement Learning* makes use of indirect feedback typically based on numerical rewards given during the interaction, and the goal is to maximise the rewards in the long run. The environment of a reinforcement learning agent is represented with a Markov Decision Process (MDP) or a generalisation of it. Its solution is a policy that represents a weighted mapping from states (situations that describe the world) to verbal and/or physical actions, and can be found through a trial and error search in which the agent explores different action strategies in order to select the one with the highest payoff. This framework can be seen as a very weak form of supervised learning, where the impact of actions is rated according to the overall goal (e.g. fetching and delivering an object or playing a game). This form of learning has been applied to design the dialogue strategies of interactive robots using MDPs [12], Semi-MDP to scale up to larger domains [7], and Partially Observable MDPs to address interaction under uncertainty [20, 19].
- *Unsupervised learning* addresses the challenge of learning from unlabelled data. Since it does not receive any form of feedback, it has to find patterns in the data solely based on its observable features. The task of an unsupervised learning algorithm is thus to uncover hidden structure in unlabelled data. This form of machine learning has been used by [19] to cluster the observation space of a POMDP-based dialogue manager, by [12] to cluster social states for multiparty interaction, and by [16] to select features for gesture recognition tasks.
- *Active learning* includes a human directly within the learning procedure assuming three data sets: a small set of labelled examples, a large set of unlabelled examples, and chosen examples. The latter are built in an interactive fashion by an active learning algorithm who queries a human annotator for labels it is most uncertain of. This form of learning has been applied to *learning from demonstration* scenarios by [2] and closely related by [15, 21].

Other forms of machine learning that can be applied to conversational robots include transfer and multi-task learning, lifelong learning, and multiagent learning, among others [8, 4]. Furthermore, while a single form of learning can be incorporated into conversational robots, combining multiple forms of machine learning can be used to address perception, action and communication in a unified way. The next section describes some challenges that require further research for the advancement of intelligent conversational robots.

3 Challenges: What is missing?

Table 1 shows a list of binary features for the robots described above. These features are grouped according to language, robotics, learning, and evaluation. The lowest numbers in the last column indicate the dimensions that have received little attention. From this table, it can observed that the main demand to be addressed is conversational robots that interact with real people in uncontrolled environments rather than recruited participants in the lab. The research directions demanding further attention are briefly described as follows:

- **Noise and crowds:** most (if not all) interactive robots have been trained and tested in lab or controlled conditions, where no noise or low levels of noise are exhibited–see Table 1. A future direction concerning the whole multidisciplinary community lies in training and evaluating interactive robots in environments including people with real needs. This entails dealing with dynamic and varying levels of noise (from low to high), crowded environments on the move, distant speech recognition and understanding [26, 23] possibly combined with other modalities [5], and real users from the general population rather than just recruited participants.
- **Unknown words and meanings:** most interactive robots have been equipped with static vocabularies and lack grammar learning (see line 5 in Table 1), where the presence of unseen words lead to misunderstandings. Equipping robots with mechanisms to deal with the unknown could potentially make them more usable in the real world. This not only involves language understanding but also language generation applied to situated domains [9].
- **Fluent and flexible interaction:** when a robot is equipped with verbal skills, it typically uses a rigid turn-taking strategy and a predefined dialogue flow (see line 6 in Table 1). Equipping robots with more flexible turn-taking and dialogue strategies, so that people can say or do anything at any time, would contribute towards more fluent and natural interactions with humans [7].
- **Common sense spatial awareness:** most conversational robots have been equipped with little awareness of the dynamic entities and their relationships in the physical world (see lines 13 and 16 in Table 1). When a robot is deployed in the wild, it should be equipped with basic spatial skills to plan its verbal and non-verbal behaviour. In this way, spatial representations and reasoning skills may not only contribute to safe human-robot interactions but also with opportunities to exhibit more socially-acceptable behaviour. See [22, 10] for detailed surveys on social interactive robots.
- **Effective and efficient learning from interaction:** interactive robots are typically trained in simulated or controlled conditions. If a robot is to interact in the wild, it should be trained with such kinds of data. Unfortunately, that is not enough because moving beyond controlled conditions opens up multiple challenges in the way we train interactive robots such as the following:
  - robot learning from unlabelled or partially labelled multimodal data (see lines 21 and 23 in Table 1) should produce safe and reasonable behaviours;
  - altering the robot's behaviour, even slightly, should be straightforward rather than requiring a substantial amount of human intervention (e.g. programming);
  - inducing robot behaviours should exploit past experiences from other domains rather than inducing them from scratch; and
  - learning to be usable and/or accepted by people from the general population is perhaps the biggest challenge.

4 Conclusion

Previous work has shown the increase in multidisciplinary work to realise intelligent conversational robots. Although several challenges remain to be addressed by specialised communities, addressing them as a whole is the end-to-end challenge that sooner or later it has to be faced. This challenge involves two crucial actions with little attention so far (a) to bring robots out of the lab to public environments, and (b) to demonstrate that they are usable and accepted by people from the general public. We hope that the topics above will encourage further multidisciplinary discussions and collaborations.

REFERENCES

[1] Dan Bohus, Chit W. Saw, and Eric Horvitz, 'Directions robot: In-the-wild experiences and lessons learned', in *AAMAS*, (2014).

[2] Maya Cakmak and Andrea Lockerd Thomaz, 'Designing robot learners that ask good questions', in *HRI*, (2012).

[3] Sonia Chernova and Andrea Lockerd Thomaz, *Robot Learning from Human Teachers*, Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers, 2014.

[4] Heriberto Cuayáhuitl and Nina Dethlefs, 'Dialogue systems using online learning: Beyond empirical methods', in *NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Comm.*, (2012).

[5] Heriberto Cuayáhuitl, Lutz Frommberger, Nina Dethlefs, Antoine Raux, Matthew Marge, and Hendrik Zender, 'Introduction to the special issue on machine learning for multiple modalities in interactive systems and robots', *TiiS*, **4**(3), (2014).

[6] Heriberto Cuayáhuitl and Ivana Kruijff-Korbayová, 'An interactive humanoid robot exhibiting flexible sub-dialogues', in *NAACL-HLT*, (2012).

[7] Heriberto Cuayáhuitl, Ivana Kruijff-Korbayová, and Nina Dethlefs, 'Nonstrict hierarchical reinforcement learning for interactive systems and robots', *TiiS*, **4**(3), 15, (2014).

[8] Heriberto Cuayáhuitl, Martijn van Otterlo, Nina Dethlefs, and Lutz Frommberger, 'Machine learning for interactive systems and robots: A brief introduction', in *MLIS*. ACM ICPS, (2013).

[9] Nina Dethlefs and Heriberto Cuayáhuitl, 'Hierarchical reinforcement learning for situated natural language generation', *Natural Language Engineering*, **FirstView**, (12 2014).

[10] Terrence Fong, Illah R. Nourbakhsh, and Kerstin Dautenhahn, 'A survey of socially interactive robots', *Robotics and Autonomous Systems*, **42**(3-4), 143–166, (2003).

[11] Naoto Iwahashi, 'Robots that learn language: Developmental approach to human-machine conversations', in *International Workshop on the Emergence and Evolution of Linguistic Communication, EELC*, (2006).

[12] Simon Keizer, Mary Ellen Foster, Zhuoran Wang, and Oliver Lemon, 'Machine learning for social multiparty human-robot interaction', *TiiS*, **4**(3), 14, (2014).

[13] V Klingspor, Y Demiris, and M Kaiser, 'Human robot communication and machine learning', *Applied A.I.*, **11**, 719–746, (1997).

[14] Jens Kober, J. Andrew Bagnell, and Jan Peters, 'Reinforcement learning in robotics: A survey', *I. J. Robotic Res.*, **32**(11), 1238–1274, (2013).

[15] Evan Krause, Michael Zillich, Thomas Williams, and Matthias Scheutz, 'Learning to recognize novel objects in one shot through human-robot interactions in natural language dialogues', in *AAAI*, (2014).

[16] Cynthia Matuszek, Liefeng Bo, Luke Zettlemoyer, and Dieter Fox, 'Learning from unscripted deictic gesture and language for human-robot interactions', in *AAAI*, (2014).

[17] Michael Montemerlo, Joelle Pineau, Nicholas Roy, Sebastian Thrun, and Vandi Verma, 'Experiences with a mobile robotic guide for the elderly', in *AAAI*, (2002).

[18] Raymond J. Mooney, 'Learning to connect language and perception', in *AAAI*, (2008).

[19] Joelle Pineau and Amin Atrash, 'Smartwheeler: A robotic wheelchair test-bed for investigating new models of human-robot interaction', in *AAAI Spring Symposium on Socially Assistive Robotics*, (2007).

[20] Nicholas Roy, Joelle Pineau, and Sebastian Thrun, 'Spoken dialogue management using probabilistic reasoning', in *ACL*, (2000).

[21] Max Schwarz, Jörg Stückler, David Droeschel, Kathrin Gräve, Dirk Holz, Michael Schreiber, and Sven Behnke, 'Nimbro@home 2014 team description', in *RoboCup @ Home League*, (2014).

[22] Luc Steels, 'Social learning and verbal communication with humanoid robots', in *Humanoids*, (2001).

[23] Pawel Swietojanski, Arnab Ghoshal, and Steve Renals, 'Convolutional neural networks for distant speech recognition', *IEEE Signal Process. Lett.*, **21**(9), 1120–1124, (2014).

[24] Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, and Nicholas Roy, 'Asking for help using inverse semantics', in *RSS*, (July 2014).

[25] Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R. Walter, Ashis Gopal Banerjee, Seth J. Teller, and Nicholas Roy, 'Understanding natural language commands for robotic navigation and mobile manipulation', in *AAAI*, (2011).

[26] Jean-Marc Valin, François Michaud, Jean Rouat, and Dominic Létourneau, 'Robust sound source localization using a microphone array on a mobile robot', in *IROS*, (2003).