A System For Robot Concept Learning Through Situated Dialogue

Benjamin Kane*
University of Rochester
bkane2@ur.rochester.edu

Felix Gervits
DEVCOM Army Research Laboratory
felix.gervits.civ@army.mil

Matthias Scheutz
Tufts University
matthias.scheutz@tufts.edu

Matthew Marge
DEVCOM Army Research Laboratory
matthew.r.marge.civ@army.mil

Abstract

Robots operating in unexplored environments with human teammates will need to learn unknown concepts on the fly. To this end, we demonstrate a novel system that combines a computational model of question generation with a cognitive robotic architecture. The model supports dynamic production of back-and-forth dialogue for concept learning given observations of an environment, while the architecture supports symbolic reasoning, action representation, one-shot learning and other capabilities for situated interaction. The system is able to learn about new concepts including objects, locations, and actions, using an underlying approach that is generalizable and scalable. We evaluate the system by comparing learning efficiency to a human baseline in a collaborative reference resolution task and show that the system is effective and efficient in learning new concepts, and that it can informatively generate explanations about its behavior.

1 Introduction

An autonomous robot interacting with a human teammate will often encounter concepts[1] that are unfamiliar to it. This is especially problematic when exploring new environments where training data is naturally limited or expensive to collect and label. For example, a robot for space exploration may encounter objects or need to carry out procedures that do not exist on Earth, and for which data is scarce. A natural approach to online learning in many environments is to take advantage of the human-robot interface, enabling the robot to engage in multi-turn dialogue where it asks the human questions to learn about a novel concept.

Many of the foundations of this approach – e.g., robot architectures capable of engaging in natural language dialogue, and mechanisms for conversational grounding and question generation – have previously been explored, but were designed as piecemeal contributions, leaving a gap in the overall problem of learning concepts through dialogue.

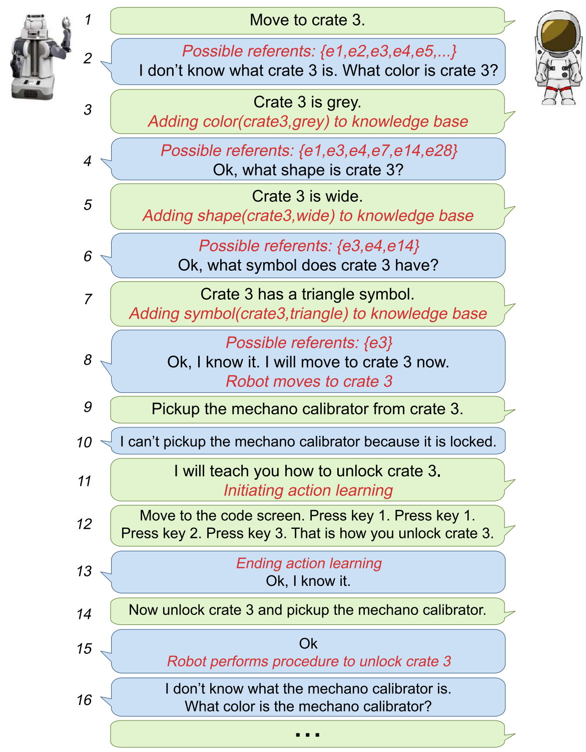

Figure 1: Example dialogue between a human and our system situated in an unexplored spacecraft environment, where the robot must learn new locations, objects, and actions through interaction with the human. The system's behavior is indicated in red.

In this work, we demonstrate a generalizable cognitive robotic system that is able to efficiently learn about unknown concepts through interactive natural language dialogue. This system leverages a probabilistic decision network model[2] to dynam-

*Work performed during a summer position at the Army Research Laboratory.

[1]The term 'concept' in this paper refers to any entity in the task domain, including objects, locations, and actions.

[2]Our decision network model, as well as the HuRDL dataset used to evaluate the system in Section 4, can be found at the following URL: https://github.com/USArmyResearchLab/ARL-HuRDL

Proceedings of the SIGdial 2022 Conference, pages 659–662
Heriot-Watt University, Edinburgh, UK. 07-09, September, 2022 ©2022 Association for Computational Linguistics

ically generate and ask optimal questions for concept learning within any environment, while also employing natural language capabilities and an explicit knowledge representation enabled by a cognitive robotic architecture. An example dialogue from our system is shown in Figure 1.

2 Background

Early work in robot concept learning through dialogue explored the use of pre-specified ontologies or graphical models to allow an agent to ask questions about objects in an environment (Lemaignan et al., 2012; Chai et al., 2018; Perera et al., 2018), or to learn actions through dialogue (She et al., 2014). Other work explores the use of proactive symbol grounding or pragmatic models for reference resolution (Williams et al., 2019; Arkin et al., 2020). In contrast to these studies, our work includes a notion of uncertainty and can scale to new task domains through dynamic adaptation of a decision network.

Recent work has built upon these approaches by introducing information-theoretic measures for selecting optimal questions. Skočaj et al. (2011) propose a robot that can ask questions about object properties that maximize information gain, and test the system using colors and shapes as properties. Deits et al. (2013) relatedly demonstrate a system that can instantiate templatic questions to minimize entropy of the robot's probabilistic symbol grounding function. Both approaches, however, rely on the use of a small fixed set of properties or question templates; we present a scalable approach that can generate questions from arbitrary properties.

3 System Design

Our system combines a decision network model for question selection (Gervits et al., 2021a) with the DIARC (Distributed Integrated Affect Reflection Cognition) robotic architecture (Scheutz et al., 2019) in order to enable interactive concept learning. The DIARC architecture, which follows a distributed, component-based design, allows for semantic parsing, introspection on knowledge, explanation generation, and support for one-shot learning of actions. The particular configuration of DIARC used by our system is shown in Figure 2. In the remainder of this section, we describe the primary components of this architecture.

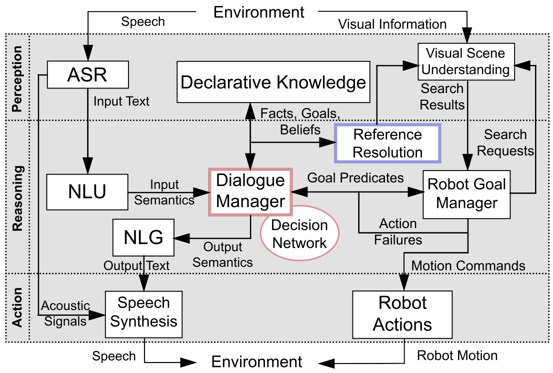

Figure 2: Architecture of the system's DIARC configuration. The core components that drive concept learning are the dialogue manager, which interacts with a decision network for question generation, and a reference resolution component for resolving concepts in user instructions to observed objects in the environment.

3.1 Decision Network

The dialogue manager component of our DIARC configuration is extended with a decision network model (Gervits et al., 2021a) that combines a Bayesian network with action and utility nodes. The model represents the robot's knowledge for a target referent and selects a question to help reduce ambiguity and acquire new concept knowledge.

Figure 3 shows a generic example of a decision network constructed by the system. The green boxes represent *chance nodes* which are random variables corresponding to the agent's knowledge of the object properties, the number of target referents, and the instruction. The blue diamond is a *utility node* which represents the utilities associated with asking questions from the red *decision node* conditioned on the chance nodes.

Since the robot's goal in asking a question is to reduce ambiguity (in the case of reference resolution, narrowing down the number of possible referents for a concept), the model selects a "best" question by calculating maximum expected utility from the model, with utilities set by calculating the Shannon entropy for each object property.

As shown by Gervits et al. (2021a), this approach is well-suited to dialogue learning in novel environments because the decision network is dynamically constructed for any novel environment given only observed object properties. Moreover, the network is constructed with the minimum set of nodes needed to disambiguate all entities in the environment, and can be re-constructed on the fly if new entities are discovered. This greatly enhances the

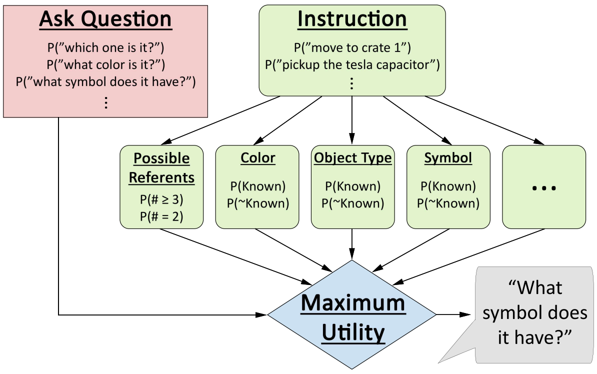

Figure 3: An instance of the decision network produced in our evaluation domain. The probabilistic chance nodes are shown in green. The red node represents the decisions available to the system, while the blue node represents the utilities associated with each decision, and outputs the decision with maximum expected utility.

flexibility of the approach, enabling it to generalize and scale to a variety of unexplored environments.

3.1.1 Semantic Parser and Declarative Knowledge

The NLU component uses a CCG grammar to map input text to a logical semantic representation[3], including the speech act type of the input (e.g., instruction or statement). The system is also able to use pragmatic inference rules to further reason about the contextual meaning of the user's utterance. The system maintains a declarative knowledge base of the system's beliefs, such as observed properties of objects, interpretations from the NLU component, and any logical inferences thereof.

3.1.2 Goal-based Dialogue Manager and Robot Actions

The dialogue manager component is responsible for handling the semantics of a speaker's input and forming system goals based on the speech act type of the user's input. In the case of an instruction, the intent of the speaker will be adopted as the robot's goal, which will either be handled by invoking an action satisfying the goal (if all referents are known), or using the decision network to generate a clarification question. In the case of a statement, the system will modify its declarative knowledge with any facts expressed in (or inferred from) the input. In both cases, the NLG component will be used to create a response by the robot; typically a simple acknowledgement.

Robot actions are implemented as *action scripts* that provide abstract logical formulations of actions consisting of preconditions, effects, and constituent steps (Scheutz et al., 2019). In our system, the robot has action scripts for every basic action that it is able to perform, such as moving to a location or picking up an object. Furthermore, DIARC allows for one-shot learning of novel actions through issuing sequences of lower-level instructions (Scheutz et al., 2017).

3.1.3 Reference Resolution

Our system is able to learn novel objects through a reference resolution component that interacts with the dialogue manager. When an unknown referent is encountered, the system will compute the number of possible entities that it could refer to, based on the properties that the system currently knows about the concept. If there are multiple possible referents, the dialogue manager will utilize the decision network model to generate a clarification question; any responses from the user are interpreted and used to update the system's declarative knowledge. Once a single referent is obtained, the system will identify the object with the corresponding concept and execute the instruction. Thus, the system is able to acquire knowledge about concepts through repeated application of this process.

4 Evaluation

To evaluate the integrated system, we implemented it on a PR2 robot in a virtual spacecraft environment containing unknown objects and procedures for the robot to learn. The robot performed a collaborative tool organization task in which it was instructed via typed natural language commands to place novel tools in their correct containers. In our evaluation, the robot is given sequences of commands from a subset of the Human-Robot Dialogue Learning (HuRDL) corpus (Gervits et al., 2021b) consisting of dialogues from 10 participants[4]. The human-generated questions in these dialogues are compared to the questions generated by the robot for the same commands in terms of *accuracy* (the proportion of commands that the robot is able to execute after resolving unknown referents) and *question efficiency* (the average number of questions that the agent must ask to learn each new concept).

The spacecraft environment contains 18 tools, with six main types and three instances of each type

[3]The logical representation used by DIARC is an extension of first-order predicate logic (Scheutz et al., 2019).

[4]We use only "low-level" dialogues with Commander initiative from the HuRDL corpus to match the robot task.

Table 1: Comparison of human performance to integrated system on question efficiency and accuracy.

	Human (N=10)	Robot (N=1)
# Questions	31	55
Question Ef.	1.72	2.29
Accuracy	0.77	1.00

(given novel sci-fi names, such as “electro capacitor”) that also vary along six feature dimensions such as color, size, etc. The environment also contains 18 containers such as platforms, lockers, and crates; some of these are locked and require learning specialized procedures to open. The robot starts with a basic perceptual representation of the entities in the environment, including their observed properties (e.g., an entity is red, small, etc.), but without a name for any of them.

Our results are summarized in Table 1[5]. Overall, the robot asked more questions than the humans on average, but attains a higher accuracy, being able to resolve every entity in the task with enough questions. These results highlight a trade-off between accuracy and question efficiency relative to human performance: as our system lacks commonsense knowledge that humans are able to draw upon when learning new concepts, it generally needs to ask more questions per object, but its systematic approach to disambiguation allows it to avoid mistakes that humans would occasionally make, such as overlooking an entity in the environment.

5 Conclusion and Future Work

We presented a robotic system that combines a decision network model for question generation with a cognitive robotic architecture to allow the system to efficiently learn about new concepts in unexplored environments through dialogue. The design of our system is scalable due to the dynamic construction of the decision network, while the robotic architecture allows for broader situated interaction including symbolic reasoning and explanation generation. Our evaluation demonstrated that our system, while having slightly lower question efficiency than human participants on the same task, was adept at learning new concepts in our experimental setting. In the future, we aim to allow the robot to automatically acquire property knowledge through exploration prior to concept learning.

[5]Since the robot produces deterministic outcomes for the same command, we perform only a single trial for the robot.

References

J. Arkin, R. Paul, S. Roy, D. Park, N. Roy, and T. M. Howard. 2020. Real-time human-robot communication for manipulation tasks in partially observed environments. In *Proc. of ISER*.

Joyce Y. Chai, Qiaozi Gao, Lanbo She, Shaohua Yang, Sari Saba-Sadiya, and Guangyue Xu. 2018. Language to action: Towards interactive task learning with physical agents. In *Proc. of IJCAI*.

Robin Deits, Stefanie Tellex, Pratiksha Thaker, Dimitar Simeonov, Thomas Kollar, and Nicholas Roy. 2013. Clarifying commands with information-theoretic human-robot dialog. *JHRI*, 2(2).

Felix Gervits, Gordon Briggs, Antonio Roque, Genki A. Kadomatsu, Thurston Dean, Matthias Scheutz, and Matthew Marge. 2021a. Decision-theoretic question generation for situated reference resolution: An empirical study and computational model. In *Proc. of ICMI*.

Felix Gervits, Antonio Roque, Gordon Briggs, Matthias Scheutz, and Matthew Marge. 2021b. How should agents ask questions for situated learning? An annotated dialogue corpus. In *Proc. of SIGdial*.

Séverin Lemaignan, Raquel Ros, E Akin Sisbot, Rachid Alami, and Michael Beetz. 2012. Grounding the interaction: Anchoring situated discourse in everyday human-robot interaction. *International Journal of Social Robotics*, 4(2).

Ian Perera, James Allen, Choh Man Teng, and Lucian Galescu. 2018. Building and learning structures in a situated blocks world through deep language understanding. In *Proc. of SpLU*.

Matthias Scheutz, Evan Krause, Brad Oosterveld, Tyler Frasca, and Robert Platt. 2017. Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture. In *Proc. of AAMAS*.

Matthias Scheutz, Thomas Williams, Evan Krause, Brley Oosterveld, Vasanth Sarathy, and Tyler Frasca. 2019. An overview of the distributed integrated affect and reflection cognitive DIARC architecture. In *Cognitive Architectures*.

Lanbo She, Shaohua Yang, Yu Cheng, Yunyi Jia, Joyce Chai, and Ning Xi. 2014. Back to the blocks world: Learning new actions through situated human-robot dialogue. In *Proc. of SIGdial*.

Danijel Skočaj, Matej Kristan, Alen Vrečko, Marko Mahnič, Miroslav Janíček, Geert-Jan M. Kruijff, Marc Hanheide, Nick Hawes, Thomas Keller, Michael Zillich, and Kai Zhou. 2011. A system for interactive learning in dialogue with a tutor. In *Proc. of IROS*.

Tom Williams, Fereshta Yazdani, Prasanth Suresh, Matthias Scheutz, and Michael Beetz. 2019. Dempster-shafer theoretic resolution of referential ambiguity. *Autonomous Robots*, 43(2).