Learning Grounded Communicative Intent from Human-Robot Dialog

Joseph Modayil
Reinforcement Learning and Artificial Intelligence Laboratory
Department of Computing Science, University of Alberta

Abstract

Studying how a robot can learn to communicate with a person provides insight into how communication might be learned in general. Deep models of dialog and communicative intent typically rely on modeling the internal state of the speakers—states that are unobservable by a learning robot. This paper considers how communication can be framed to be learnable from experience. In particular, we describe how an agent might learn to communicate by building on three foundational capabilities, namely 1) an observable signal of satisfied intent (a smile), 2) the ability to imitate perceived actions, and 3) perceptual referents for discourse items. Early simulation results show that an agent can learn some basic communication skills from these foundations.

Introduction

This paper explores the premise that a robot might be able to efficiently learn to communicate with cooperative people by building on three foundational competencies: 1) a shared signal of satisfied communicative intent, 2) the ability to imitate observed actions, and 3) shared perceptual references for discourse items. By combining these abilities, a robot might be able to bridge the divide between robot skills and the ability to communicative effectively with people.

Much work in dialog and communication is founded on modelling the beliefs and intents of speakers (Cohen and Levesque 1990; Kautz and Allen 1986). Semantics of speech acts in these models comes from people manually binding grammatical rules with human constructed logical forms. Although these intention models have been successfully deployed on interactive systems (Allen et al. 2007) and robot systems (Roy, Pineau, and Thrun 2000), they rely on manually constructed models of grammar and state.

Another line of research looks at how communication and shared intents arise in simple communication systems. This includes philosophical examination of communication as a cooperative game (Lewis 1969), robot studies into how a primitive language can arise between cooperative agents (Steels 2005), and biological studies into the mental processes that support communication (Tomasello 2003).

Grounding communicative intent in experience can bridge these two lines of research.

Proposal

Several problems in learning the semantics of communication might be overcome by combining three functional capabilities on a robot: a signal of satisfied intent, imitation, and perceptual referents.We present an extended example that describes the scenario that generates the dialog in Figure 1. Assume that there is a learning robot (the student) interacting with a cooperative person (the teacher).

1. The teacher points and asks for the name of a fruit.
2. The student is unsure of what to do and babbles.
3. The teacher demonstrates the desired response.
4. The student imitates the teacher.
5. The teacher smiles, indicating a satisfied intent.

By acting randomly, with enough time the robot could learn from experience what action it can perform that will likely make the person smile. However if the robot engages in imitation when it is unsure of the desired response, the interaction will provide relevant experience. Generalizing from this example involves the three capabilities mentioned above.

1. *Signalling Satisfied Intent-* If a person asks a robot to point to an apple and eventually the robot does, then the person's smile can serve as an observable signal to the robot that it has satisfied the person's intent. Instead of reasoning directly with the unobservable beliefs of the person, the robot can act so as to maximize the probability that the person will smile. Using this observed signal, the robot's experience will contain sufficient information to indicate a successful interaction, independent of the robot's model of the person's state.
2. *Imitation-* Predicting that a performed action will make a requester smile is not the same as actually being able to perform the action. The second requirement is the ability to imitate an observed action. There can be multiple imitations that are not useful, and imitation capabilities will vary with the robot's knowledge. However, even a limited imitation ability can facilitate knowledge acquisition.
3. *Perceptual Referents-* Without prior linguistic knowledge or perceptual referents, a robot can not generalize across dialogs. With perceptual referents, a robot can generalize across attributes, for example recognizing that pointing at one location with an apple is similar to pointing at another location with an apple. The robot can also learn associations between perceived attributes (such as shape) and

Per(0, :tau801, <shape>, :Orange), Per(0, :tau801, <col>, :Green), Per(0, :tau801, <size>, :Small), Per(0, :tau801, <loc>, :(1, 3)), // perceptual referents
Per(0, :tau802, <shape>, :Apple), Per(0, :tau802, <col>, :Red), Per(0, :tau802, <size>, :Small), Per(0, :tau802, <loc>, :(1, 4)),
Act(1, :teacher, <Point>, :(1, 4)), Act(1, :teacher, <Say>, :Pronounce(SayName)),
Act(2, :student, <Say>, :Pronounce(babbling)),
Act(3, :teacher, <Say>, :Pronounce(Apple)),
Act(4, :student, <Say>, :Pronounce(Apple)), // imitation
Act(5, :teacher, <Smile>, -) // satisfied intent

Figure 1: A sample training dialog represented as a set of tuples. There are tuples for (Per)ceptions and (Act)ions, whose components are time, object/agent identifier, attribute name, and attribute value. Learning is facilitated by 1) a smile indicating a satisfied intent, 2) imitating an intended response, and 3) perceptual referents.

snippets of speech (a pronunciation of apple) and generalize appropriately.

Simulation

The proposed scenario has been explored in a symbolically-driven small simulation. First a collection of dialogs were generated between a teacher and an ignorant student represented with tuples as shown in Figure 1. Each dialog has two objects, each in one of four locations. The objects come in three sizes (big, medium, small), three shapes (apple, orange, banana), and a shape appropriate color (green, yellow, red, orange). The teacher requests the name of an object attribute, or requests that the student point to some location.

Two rule learning algorithms were deployed to examine if a robot could extract rules for the intended action from this data, and the learned rules were evaluated on separate test data. The first algorithm operates in a FOIL-like manner (Quinlan and Cameron-Jones 1993)), greedily adding propositional rules to cover positive examples by finding a common subset of tuples in the dialogs with the same ground intent (the robot action preceding a smile). The second algorithm learns parametrized rules that generalize the propositional rules across parameter values and creates named associations (for example between the percept for the shape of an apple and the pronunciation of the word apple). The results from a five-fold cross-evaluation are shown in Figure 2. The results show that the system is able to learn both kinds of rules, but learns more efficiently with parametrized rules.

Discussion

This work suggests that simple conventions in natural communication enable an agent to efficiently learn intent-driven communication rules. This simulation shows the potential for a robot to learn communicative intent directly from experience, using an observable signal of satisfied intent, the ability to imitate actions, and perceptual referents for the dialog. The simulation suggests generalization is better with

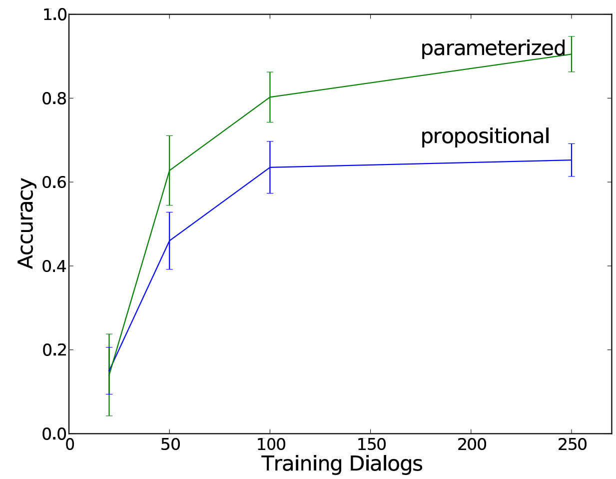

Figure 2: Accuracy of learned responses from a greedy propositional rule learner and parametrized rule learner. Error bars show standard deviation on 5-fold cross validation.

more expressive representations, but this step can be considered an optimization. The presented method builds on skills that have been demonstrated on robots platforms: perceiving objects, learning multi-modal associations, establishing joint attention, learning from demonstration and imitating actions (Modayil and Kuipers 2007; Roy and Pentland 2002; Nehaniv and Dautenhahn 2007). Further work is required to generalize this approach to richer communication (real settings, compositional rules, and multi-step dialogs).

References

Allen, J.; Chambers, N.; Ferguson, G.; Galescu, L.; Jung, H.; Swift, M.; and Taysom, W. 2007. PLOW: A collaborative task learning agent. In *Proc. 22nd Nat. Conf. on Artificial Intelligence*.

Cohen, P., and Levesque, H. 1990. Intention is choice with commitment. *Artificial Intelligence* 42:213–261.

Kautz, H., and Allen, J. 1986. Generalized plan recognition. In *Proc. 5th Nat. Conf. on Artificial Intelligence*, 32–37.

Lewis, D. 1969. *Convention: A Philosophical Study*. Harvard University Press.

Modayil, J., and Kuipers, B. 2007. Autonomous development of a grounded object ontology by a learning robot. In *Proc. 22nd Nat. Conf. on Artificial Intelligence*.

Nehaniv, C., and Dautenhahn, K., eds. 2007. *Imitation and Social Learning in Robots, Humans and Animals*. Cambridge University Press.

Quinlan, J. R., and Cameron-Jones, R. M. 1993. FOIL: A midterm report. In *Proceedings of the European Conference on Machine Learning*, 3–20. Springer-Verlag.

Roy, D. K., and Pentland, A. P. 2002. Learning words from sights and sounds: a computational model. *Cogn. Science* 26:113–146.

Roy, N.; Pineau, J.; and Thrun, S. 2000. Spoken dialogue management using probabilistic reasoning. In *Proc. of the 38th Ann. Meeting of the Assoc. for Computational Linguistics (ACL-2000)*.

Steels, L. 2005. The emergence and evolution of linguistic structure: from lexical to grammatical communication systems. *Connection Science* 17(3-4):213–230.

Tomasello, M. 2003. *Constructing a Language*. Harvard University Press.