Incremental Knowledge Acquisition with Selective Active Learning

(Doctoral Consortium)

Batbold Myagmarjav
Supervisor: Mohan Sridharan
Texas Tech University
Lubbock, Texas 79409, USA
bat.myagmarjav@ttu.edu

ABSTRACT

This paper describes an architecture for robots interacting with non-expert humans to incrementally acquire domain knowledge. Candidate questions are generated using contextual information and ranked using different measures with the objective of maximizing the potential utility of the response. We report results of some preliminary experiments evaluating the architecture in a simulated environment.

Categories and Subject Descriptors

I.2.6 [**Learning**]: Knowledge Acquisition

General Terms

Human Factors, Algorithms

Keywords

Human-robot collaboration, Knowledge Acquisition, contextual query generation

1. INTRODUCTION

Effective Human-robot collaboration in complex domains typically requires considerable domain knowledge and a large number of labeled samples of interesting objects and events. However, humans may not have the time or expertise to provide elaborate and accurate information, and it may not be feasible to provide labeled examples of all objects and events of interest. Researchers have designed many active learning algorithms to allow incremental labeling or acquisition of data e.g., instances that have been misclassified using existing models [5]. Existing algorithms tend to focus on choosing instances to be labeled. Even when active learning is combined with multiple instance learning to minimize human supervision [6], the algorithms do not pay much attention to the types of queries. Research indicates that the queries that allow labeling of features and object instances significantly improve knowledge acquisition [3]. AI researchers have developed algorithms for agents to embed context to improve the quality of the questions being posed in a human-robot interaction (HRI) setting [4], but such algorithms tend to focus on the human reaction and ability to answer these questions. Posing query generation as a planning task is also challenging because it requires prior knowledge of possible answers, which may be different for different scenes, and is likely to be computationally inefficient. More recent research has combined active learning with learning from demonstration to explore different types of questions [1], but has focused on how query categories are perceived by humans.

Success of active learning in HRI thus requires an effective strategy for posing questions that supports faster learning based on limited interaction with humans who may not have domain expertise. Towards this objective, our architecture allows a robot (i.e., learner) to use contextual cues to generate candidate queries. From these candidates, the robot incrementally poses questions that have high relative utility, i.e., questions that disambiguate between, and quickly acquire information about, domain objects. We report results of some preliminary experiments evaluating the architecture in a simulated environment.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

2. METHODS

Figure 1(a) is an overview of the incremental selective active learning architecture in the context of simulated images of scenes with objects with different properties (color, shape and size). Queries are generated with differing levels of context and ranked according to their utility and contextual unambiguity as discussed in the following sections.

2.1 Utility

Humans are often not readily available in active learning settings. It is thus best to maximize the utility of the question posed to increase the chance of learning the most from the human answer. One query selection criteria would be the potential information gained about the *scene* by posing a specific question. The learner agent would keep track of frequency of a certain feature or object instances present in the *scene*. Acquiring the label of the object or the feature instance that is most frequently encountered would maximize the utility of limited interaction with a human annotator.

2.2 Context and Unambiguity

The ability to pose relevant questions that quickly draw a human's attention to the object(s) of interest can significantly influence the quality of a robot's interaction with

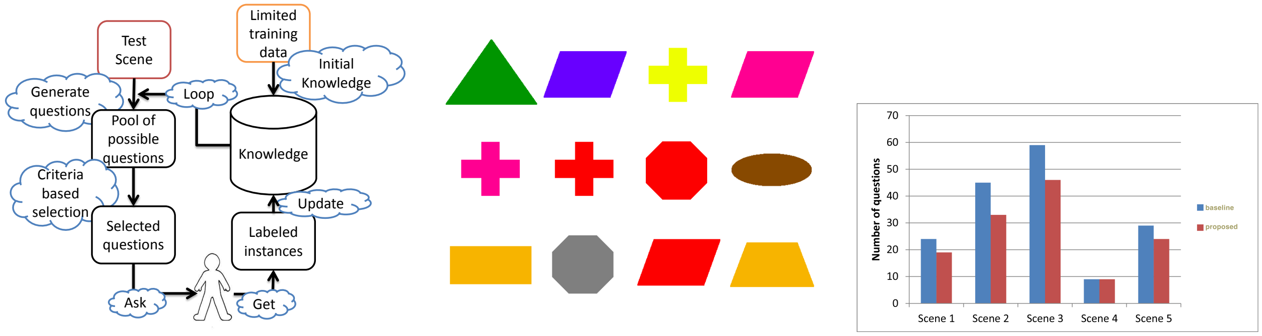

(a) Architecture (b) Illustrative scene (c) Experimental results

Figure 1: Figure 1: (a) Proposed architecture; (b) An example scene; (c) Proposed architecture requires less number of queries than baseline algorithm.

humans. Humans frequently use contextual cues to draw attention to an object of interest. Contextual cues can take different forms, and positional context with reference to a known object can be very useful in disambiguating the object of interest. For instance, instead of referring to a "1965 Ford Mustang" in a busy street intersection, we may refer to the "Red car behind the bus", using both feature labels (e.g., color and object labels) and positional reference to a known object.

2.3 Query types

We consider candidate queries based on different contextual cues, and rank them based on measures of information gain, ambiguity and human confusion. Some examples of possible queries are posed below:

- *Object Label*: "What is the label of the object in the bottom right of the scene?"
- *Feature Label*: "What is the label of the color of object in top left of the scene?"
- *Affirmation*: "Is there a blue object in the scene?"
- *Location*: "Where is a blue object?"

3. RESULTS AND DISCUSSION

We report preliminary results of evaluating our architecture in a simulated domain [2]. We abstract away the nondeterminism in object recognition and speech understanding; objects are recognized once the individual features are learned, and speech gets translated into text and parsed to generate the labels. The trials summarized below consider simulated images with objects characterized by 10 different colors and 15 different shapes. As an illustrative example of query generation, consider the scene in Figure 1(b). Assume that the color and shape labels of four objects are known a priori:*pink star, green arrow, blue heart,* and *yellow cross*. The following are some of the queries generated; each line ends with the answer provided to the question:

- *Iteration 4*: "What is the label of the object in the bottom right of the scene?" **Orange Trapezoid**.
- *Iteration 6*: "What is the label of the object that is to the left of the orange trapezoid?" **Red Parallel**.
- *Iteration 13*: "What is the label of the object that is above the red parallel?" **Red Octagon**.

The information obtained by posing questions is used to formulate and pose subsequent questions, and questions may refer to more than one object. Figure 1(c) summarizes results for five scenes. These scenes differ in terms of the number and type of objects in the scene. For instance, "Scene 1" corresponds to Figure 1(b), while "Scene 2" and "Scene 3" have 20 and 30 objects respectively. For each scene, the robot started with the same initial knowledge about a subset of the objects. As a baseline for comparison, we used an algorithm that started with the same initial knowledge but selected queries randomly from the set of candidate queries. We observe that the number of queries required to acquire the desired labels increases as the scene becomes more complex. However, for each set of paired experimental trials, our architecture results in the desired labels of objects and features being acquired by posing a much smaller number of queries. Similar results were obtained over 100 different (randomly generated) scenes with different number and types of objects. Future work will implement and evaluate the architecture in more complex scenes, and on physical robots interacting with non-expert humans.

REFERENCES

[1] M. Cakmak and A. Thomaz. Designing robot learners that ask good questions. In *Human-Robot Interaction (HRI), 2012 7th ACM/IEEE International Conference on*, pages 17–24, March 2012.

[2] B. Myagmarjav and M. Sridharan. Incremental knowledge acquisition with selective active learning. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 4-8, 2015.

[3] H. Raghavan, O. Madani, and R. Jones. Active learning with feedback on both features and instances. *Journal of Machine Learning Research*, 7:1655–1686, 2006.

[4] S. Rosenthal, A. K. Dey, and M. Veloso. How robots' questions affect the accuracy of the human responses. In *Robot and Human Interactive Communication, 2009. RO-MAN 2009. The 18th IEEE International Symposium on*, pages 1137–1142, 2009.

[5] B. Settles. *Active Learning*. Morgan & Claypool publishers, 2012.

[6] B. Settles, M. Craven, and S. Ray. Multiple-instance active learning. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1289–1296. Curran Associates, Inc., 2008.