Incremental Knowledge Acquisition with Selective Active Learning

(Extended Abstract)

Batbold Myagmarjav
Computer Science
Texas Tech University
Lubbock, Texas 79409 USA
bat.myagmarjav@ttu.edu

Mohan Sridharan
Electrical and Computer Engineering
The University of Auckland
Auckland 1142, New Zealand
m.sridharan@auckland.ac.nz

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

Algorithms, Human Factors

Keywords

Human-robot collaboration, knowledge acquisition, contextual query generation.

1. INTRODUCTION

Human-robot collaboration in complex domains frequently 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 [4], but these algorithms tend to focus on choosing instances to be labeled. Even when active learning is combined with other approaches such as multiple instance learning to minimize human supervision [5], there is not much focus on the types of queries. Research indicates that the queries that allow labeling of features and object instances significantly improve knowledge acquisition [2]. Algorithms have also been developed for agents to embed context to improve the quality of the questions being posed to humans [3], but such algorithms tend to focus on a human's reaction and
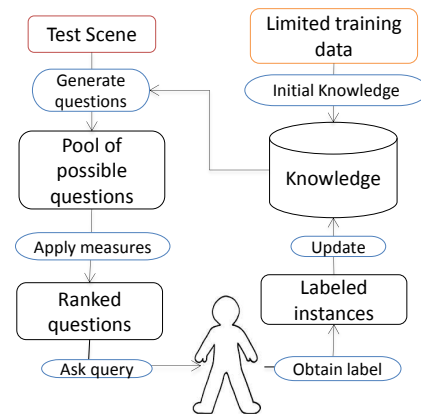
Figure 1: Proposed architecture.

ability to answer these questions. Posing query generation as a planning task is also challenging because it requires prior knowledge of the possible questions and answers, which may vary for different scenes. 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. The architecture described in this paper seeks to address these limitations by posing questions that support faster learning based on limited interaction with humans who may not have domain expertise. We report results of preliminary experiments evaluating the architecture in a simulated environment.

2. PROPOSED ARCHITECTURE

Figure 1 shows the proposed architecture that generates candidate queries using contextual cues. In this paper, we consider the relative position of an object , e.g., *above*, *below*, *left of*, *right of*) with unknown object or features labels, with respect to objects with known labels (*local context*), and with respect to the entire image (*global context*). These queries are ranked based on relative *utility*, i.e., the ability to disambiguate between, and quickly acquire information about, desired objects and events. We compute utility as a combination of measures of information gain, ambiguity, and human confusion. Top-ranked queries are used to solicit human feedback when it is available, and the revised knowledge is used to generate subsequent questions. The individual measures are summarized below.

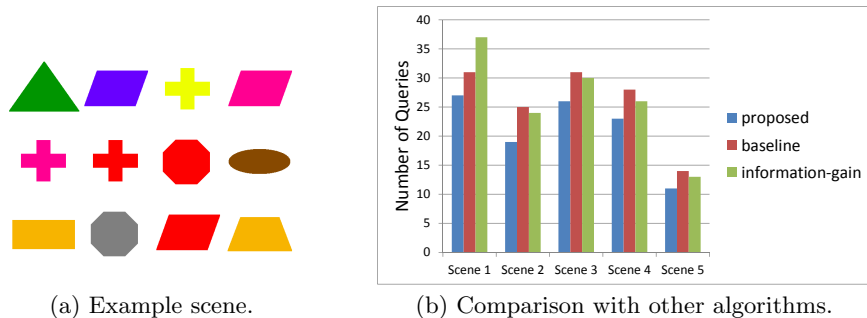(a) Example scene.　　　　(b) Comparison with other algorithms.

Figure 2: (a) Example scene; (b) Proposed architecture requires a much smaller number of queries than algorithms that select queries randomly, or use only the information gain measure.

Information Gain: The first measure (*information gain*) seeks to maximize the information that can potentially be gained by posing specific questions. It is based on the observation that acquiring the label of the feature or object instance that is most frequently encountered would maximize the utility of soliciting help from a human annotator. To compute the value of this measure for each candidate query, the system keeps track of the frequency of occurrence of each feature or object instance in the scene.

Unambiguity: The second measure (*unambiguity*) assigns a high score to a query whose embedded context uniquely identifies an object or event in the scene. This measure is modeled as a modified Chi-square distribution whose value for a specific query drops as the number of objects satisfying the context embedded in the query increases.

Human Confusion: The third measure (*human confusion*) captures the fact that humans can become confused as the amount of contextual information included in a query increases. This measure assigns a score to each query inversely proportional to the amount of context embedded in the query. The net measure of utility is the product of the information gain and unambiguity measures, with the measure of human confusion used to break ties.

3. RESULTS AND DISCUSSION

To thoroughly analyze and understand the contribution made by the proposed algorithm and each measure, we abstract away the non-determinism 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 below consider simulated images of scenes with objects characterized by one of 10 different colors and 15 different shapes; object labels are a combination of the color and shape labels.

First, consider an illustrative example of query generation in the simple scene in Figure 2(a). Assume that the color, shape, and object 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 color that is to the left of the orange trapezoid?" **Red**.

Next, we compared the proposed query selection algorithm with an algorithm that uses only the unambiguity measure, by computing the % knowledge of object and feature labels in the scene as a function of the number of queries posed to acquire this knowledge. These algorithms take the same number of queries to obtain complete knowledge of the scene, but the proposed algorithm allows the robot to maximize the knowledge acquired during each interaction with a human.

Finally, Figure 2(b) summarizes results for five scenes which differ in terms of the number and type of objects. The proposed algorithm for selecting questions from the set of candidate questions was compared with the algorithm that selected questions randomly ("baseline"), and with the algorithm that only used the information gain measure to select queries; all three algorithms start with the same initial knowledge about a small set of objects. Over a set of 100 different (randomly generated) scenes with different number and type of objects, the ratio of the average number of questions posed using just the information gain measure with the number of question posed using our algorithm is 1.19 ± 0.112; the ratio when the random selection algorithm is compared with our algorithm is 1.17 ± 0.106. These results are statistically significant, and more pronounced as the scenes become more complex. Future work will implement and evaluate the architecture in more complex scenes, and on physical robots.

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] 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.

[3] 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.

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

[5] 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.