Teaching and Learning using Semantic Labels

Sayanti Roy[1], Harshal Maske[2], Girish Chowdhary[2] and Christopher Crick[1]
[1]Oklahoma State University, Stillwater OK 74074
[2]University of Illinois at Urbana-Champaign, Champaign IL 61801
sayanti.roy@okstate.edu, hmaske2@illinois.edu, girishc@illinois.edu,
chriscrick@cs.okstate.edu

ABSTRACT

We improve robotic learning from demonstration (LfD) via an active learning process of interacting with a human expert to establish a semantic structure and labels for a sign language task. This process situates a learned task in a human-accessible conceptual framework, in order to improve skill transfer not only from expert human teacher to robot, but from robot to novice human learner.

1. INTRODUCTION

Learning from demonstration (LFD) has long been used to enable end-users to teach tasks to robots without traditional programming, hopefully increasing a robot's flexibility and availability to nonexperts. In this paper, the humanoid robot Baxter is taught a sequence of sign language motions using its left arm which communicate the sentence, "Hello, please listen to me". During the learning process, Baxter can request information from the human expert using label, feature and demonstration queries [1]. The robot is able to segment, hierarchically structure and evaluate the confidence of components of the demonstrated task, and then solicits feedback from the expert demonstrator to help this process.

The process of interactively refining and labeling a semantic task hierarchy improves the robot's learning from demonstration, as illustrated in our experimental results. Equally importantly, however, we aim to use this structure to assist in a much less-well-studied mileu, that of robotic *teaching* from demonstration. The work reported here is in service of the larger goal of creating robots with enough human-accessible task understanding to be able to act successfully as tutors and coaches for complex skill learning. Novice learners are often unable to notice important information and patterns related to a novel task, compared to experts. Hence if the entire task is decomposed into well-organized subtasks with meaningful labels [2, 7, 3], the novice should be more able to distinguish relevant relationships and will learn better.

HRI '17 March 6–9, 2017, Vienna, Austria

ACM ISBN 978-1-4503-4885-0/17/03.
DOI: http://dx.doi.org/10.1145/3029798.3038335

In this paper we quantify performance by comparing the root mean squared error (RMSE) of trajectories learned by traditional LFD and our query-driven active learning and semantic labeling. With this approach, Baxter is able to learn a complex motor task more accurately (compared to an ideal trajectory) than in the condition when the robot is merely shown a series of demonstrations of the same task.

2. TECHNICAL APPROACH

Fig. 1 shows a hierarchical task learning process using multiple action primitives. The first planning sequence is the primary goal, further decomposed into subgoals. These subgoals define series of sequential motions which comprise the task. In the figure, the tree represents a task after the learning process and active queries. The nomenclature of the subgoals is done based of the movement of a particular joint in Baxter. Blue color denotes semantic labels learned from queries and applied to subtasks. Red illustrates task segmentation primitives discovered by the robot. Green indicates primitives which were replaced with new demonstrations the during active learning process.

Demonstrations are first collected and synthesized into learned motion using dynamic time warping (DTW) and barycenter averaging [5, 6]. The demonstrated motion is then segmented to identify change points using non-Bayesian clustering (NBC) within a vector-valued Gaussian Process (VGP):

$$\begin{aligned}\log P(y|x, M) = & -\frac{1}{2}(y-\mu(X_*))^T \Sigma_{X_*X_*}(y-\mu(X_*)) \\ & - \log|\Sigma_{X_*X_*}|^{\frac{1}{2}} + C\end{aligned} \tag{1}$$

Here, $\mu(X_*)$ is the matrix-valued mean prediction and $\Sigma_{X_*X_*}$ is the conditional variance including noise [4]. This segmentation generates task primitives using change point detection, which is later built into a semantically labeled structure from demonstration queries.

3. EXPERIMENTAL RESULTS

Fig. 2 depicts the performance of Baxter during $n = 8$ teaching trials, four in the plain LfD scenario and four in the semantic labeling scenario. The error is computed by comparing the learned task with an idealized, hard-coded, perfect task execution. In the semantic labeling scenario, the robot asks for expert feedback at the motion primitive change points and requests further subtask clarification when the variance of a particular segment is high. In addition, the robot elicits semantic labels for the task hierarchy

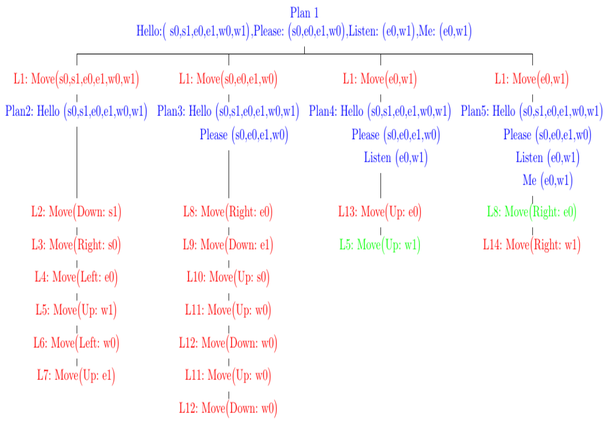

Figure 1: Hierarchical task planning using semantic labels in Baxter.

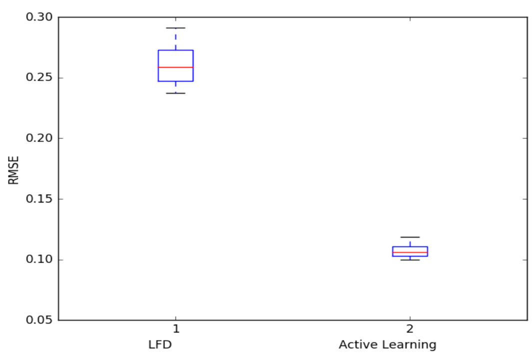

Figure 2: Root mean squared error (RMSE) for learned sign language task after demonstrations only (LFD) and after semantic label queries (active learning).

(via a text interface).

The evaluation of Baxter's task performance is conducted in an HRI setting. The initial movement demonstrations and the expert feedback during demonstration queries (for the semantic labeling scenario) are provided using a joystick controller. In the plain LfD scenario, Baxter performs a learned sign language motion generated from demonstrations given by the expert. In the semantic labeling scenario, expert feedback with suitable labels are introduced on the same learned motion to improve that sequence. The robot queries the expert for semantic labels for important motion primitives.

During active learning, demonstration queries were frequently posed to clarify ambiguous subtask segments. These disambiguation queries, in conjunction with label and hierarchy queries, improved the robot's learned model in comparison to the model developed from demonstrations alone. In the plain LfD condition, Baxter could only execute a few sections of a subtask properly, whereas the motion sequences collected after active learning were much closer to the ideal motion. Median RMSE for active learning was 0.108, while in the plain LfD condition is was 0.262 ($p < 0.002$).

4. CONCLUSION AND FUTURE WORK

We have shown that active learning and demonstration queries improve task performance in a learning from demonstration environment. Hierarchical task segmentation provides a set of building blocks which can then be interpreted and evaluated by an expert in an active learning scenario, leading to the construction of a task representation which is accessible to both humans and robots. The ultimate goal of this project is not only to increase robot performance on learned tasks, however. The learned model, constructed with semantic and domain knowledge provided by an expert teacher, should also improve the robot's ability to teach the task to a novice operator.

Future work will validate this effort. Complex machine manipulation is a skill which currently requires a long period of one-on-one human apprenticeship, for example in the control of heavy equipment in the construction industry. It is unrealistic to assume that all such machines will be replaced by completely autonomous robots in the near future, but if robots can be taught by experts while at the same time acting as coaches and tutors for novices, the overall rate of skill acquisition by both humans and robots should improve. Learned, labeled, human-oriented task organization should help this process. Our work attempts to replicate the expertise of human teachers in a robotic learning *and* teaching using expert demonstration context.

This work was supported by NSF award #1527828 (NRI: Collaborative Goal and Policy Learning from Human Operators of Construction Co-Robots).

5. REFERENCES

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

[2] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical task and motion planning in the now. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1470–1477. IEEE, 2011.

[3] G. Konidaris, S. Kuindersma, R. Grupen, and A. Barto. Robot learning from demonstration by constructing skill trees. *International Journal of Robotics Research*, 31(3):360–375, 2012.

[4] H. Maske, M. Matthews, A. Axelrod, H. Mohamadipanah, G. Chowdhary, C. Crick, and P. Pagilla. Collaborative goal and policy learning from human operators of construction co-robots. In *Neural Information Processing Systems (NIPS)*, 2014.

[5] F. Petitjean, A. Ketterlin, and P. Gançarski. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition*, 44(3):678–693, 2011.

[6] G. A. ten Holt, M. J. Reinders, and E. Hendriks. Multi-dimensional dynamic time warping for gesture recognition. In *13th Annual Conference of the Advanced School for Computing and Imaging*, volume 300, 2007.

[7] C. Wu, I. Lenz, and A. Saxena. Hierarchical semantic labeling for task-relevant rgb-d perception. In *Robotics: Science and Systems (RSS)*, 2014.