

Predicting Supportive Behaviors for Human-Robot Collaboration

Robotics Track

Elena Corina Grigore

Yale University, Department of Computer Science
New Haven, CT
elena.corina.grigore@yale.edu

Alessandro Roncone

Yale University, Department of Computer Science
New Haven, CT
alessandro.roncone@yale.edu

Olivier Mangin

Yale University, Department of Computer Science
New Haven, CT
olivier.mangin@yale.edu

Brian Scassellati

Yale University, Department of Computer Science
New Haven, CT
brian.scassellati@yale.edu

ABSTRACT

We present a model for predicting what supportive behaviors a robot should offer to a person during a human-robot collaboration (HRC) scenario. We train and test our model in simulation, using noisy data that mimics a real-world HRC interaction. Our results show that we can achieve accurate predictions, using only a small set of labeled demonstrations. We also show transfer learning capability: we train our model on an initial task and test it on a new task composed of the same building blocks but structured differently.

KEYWORDS

human-robot collaboration; robotics; learning and adaptive systems

ACM Reference Format:

Elena Corina Grigore, Olivier Mangin, Alessandro Roncone, and Brian Scassellati. 2018. Predicting Supportive Behaviors for Human-Robot Collaboration. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 3 pages.

1 INTRODUCTION

Our work tackles one of the main challenges in human-robot collaboration (HRC), that of developing robots that can learn how to be of assistance to human workers during physical tasks, such as furniture assembly. Our goal is for the robot to learn how to provide such assistance from low-level observations of human behaviors and environment features. Within this context, the robot aims to help the human worker more effectively complete the task by executing useful behaviors throughout the interaction [13].

The main challenge is that learning directly from high-frequency and high-dimensional noisy observations requires a large number of training samples. Moreover, information about the supportive behaviors to be offered requires labeled data, which is expensive to acquire in HRC. To address this, we propose a two-step approach. The robot first learns the structure of the behaviors—and hence of the task—from raw observations of human workers. The system then takes into account a few annotated demonstrations from the

human worker about what supportive behaviors are appropriate to offer throughout the task.

Our system extends hidden Markov models (HMMs) [19] in the following way. We model the task by learning from a training set consisting of observed task trajectories. These trajectories do not include information about the supportive behaviors the robot should offer. We then label a small set of these trajectories with labels for the robot supportive behaviors. We present our results on a realistic HRC task in simulation, achieving predictions with a small number of errors by using as few as three labeled trajectories.

If our robots are to adapt to novel situations, another critical challenge in robot learning is to leverage the knowledge acquired on an initial task for generalizing to similar yet unseen tasks, users, and environments. In this work, we explore this question by training the system on an initial task and testing it on a variation of the task consisting of a different structure. We hence demonstrate how our method is flexible enough to transfer to new scenarios.

Our contribution is threefold. First, we explain how to learn a model of human behaviors from ambiguous observations and to how use it to provide supportive behaviors to human workers throughout the execution of a task. Second, we demonstrate that we can do so with only a small set of provided labels. Third, we show that we can perform well on new tasks, after having learned a model on a similar but differently structured task.

2 RELATED WORK

Our work is placed between approaches for modeling and predicting human behaviors and techniques for modeling robot decisions.

Work on modeling human motions and activities includes modeling manipulation interactions [18], recognizing continuous human grasping sequences [1], and modeling multiple time series for motion capture segmentation [4]. Within robotics, we mention acquiring behavioral models for robots via HMMs [5], learning robot trajectories from demonstration [20], and robot learning to reproduce gestures by imitation [3].

Model learning in robotics encompasses work such as learning real-time inverse dynamics for robot arm control [14], modeling trajectories and motor command generation for controlling robotic manipulators [9], and learning representations for mobile robot navigation tasks that leverage interacting with people for help [15].

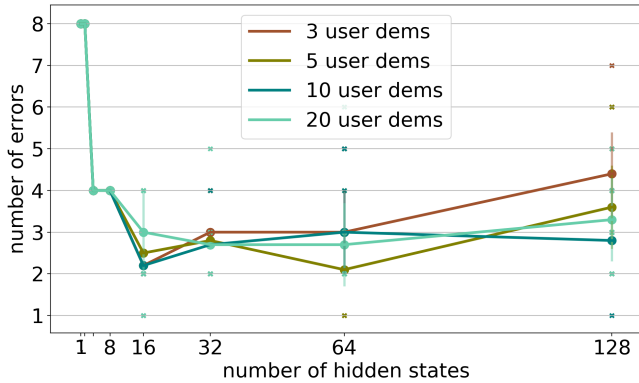


Figure 1: Number of errors as a function of the number of HMM hidden states for different sizes of trajectory labels. The number of errors is computed with 500 training trajectories and labels for 3, 5, 10, and 20 demonstrations (dems).

Work within HRC includes improving human-robot fluency [2, 7], modeling human intentions via HMMs for autonomous mobile robots [10], anticipatory control for HRC [8], predicting human motion for collaborative manipulation planning [11, 12], and learning joint action models for HRC tasks [16, 17], among others.

The originality of our method is to learn human activities as hidden variables from motion capture time series, and leverage this latent representation for providing supportive behaviors in HRC.

3 METHODS

We explore how our system predicts what supportive behaviors a robot should offer to a human worker throughout a chair assembly task [21]. We represent the task as a hierarchical task model (HTM) [6], encoding sub-states and sequencing constraints (e.g., parallel, sequential, etc.). We use the HTM only to generate the training set, and our learner does not have access to the HTM. The training set consists of trajectories that only include HTM leaves. Each leaf generates a noisy vector consisting of binary features.

To simulate collecting noisy observations in the real world, we generate a vector of $n = 5$ features for a given HTM leaf based on a signature of probability values. For example, for a leaf called *assemble leg 1*, we randomly generate a vector $[p_1, p_2, p_3, p_4, p_5]$, where $p_j \in [0; 1]$ represents the probability of feature y_t^j having value 1 at time t . We then generate the actual observations, as binary vectors, drawing from Bernoulli distributions whose parameters are given in the vector. Whenever the same leaf appears, it has the same signature of probabilities, but might have different values for its features depending on this generation process.

We then take a subset of trajectories and assign each leaf within each trajectory with a label corresponding to what supportive behavior the robot should offer. We estimate the probability for a given supportive behavior U_t at step t based on the estimated probabilities over the decoded states H :

$$P(U_t = u | H_t = k) = \frac{\sum_{i=1}^n \sum_{t=1}^T P(H_t^i = k | y_{1...T}^i) \times \mathbb{1}_{u_i=u}}{\sum_{i=1}^n \sum_{t=1}^T P(H_t^i = k | y_{1...T}^i)}.$$

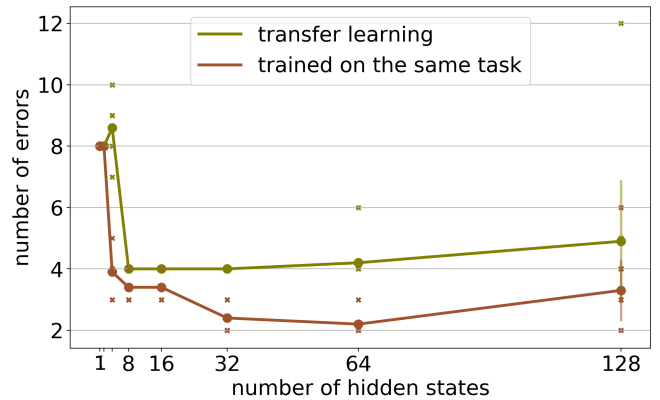


Figure 2: Transfer learning results. This graph shows the number of errors as a function of the number of HMM hidden states when tested on a transfer learning task. The model is learned on an initial task with building block A followed by B, and tested on a different task with building block B followed by A.

4 RESULTS AND CONCLUSIONS

We show the number of errors as a function of the number of HMM hidden states, averaged over ten runs for each, with increasing numbers of labeled trajectories (Figure 1). The graph highlights a trade-off between the complexity of the model and the number of trajectories we wish to label. The higher the complexity, the better the model can represent the task, yet the more labeled trajectories are necessary. In our scenario, this trade-off results in a minimum of 1 and an average number of 2 errors for a task with a total of 12 steps. Of course, this trade-off will differ for specific applications.

Another critical aspect in robot learning is that of transfer learning. When faced with a new task, a robot can speed up learning via leveraging past knowledge about similar situations. We show strong results for transfer learning, where we train on a task composed of two main building blocks (block A followed by B), and test on a different task composed of the blocks arranged differently (B followed by A). We obtain as low as 4 errors when trained entirely on the initial task (Figure 2). The baseline we compare against is training and testing on the same, initial task. For our best results, we perform worse than training on the task itself by a difference of only 1 error, which happens when the HMM has 16 hidden states.

This transfer learning result represents a success for HRC scenarios like the one we consider herein. What this amounts to in a real-world environment like a factory is being able to simply collect observations of particular sub-tasks of larger, ongoing tasks. We then only require users to label a single combination or way of arranging these sub-tasks together, and we are able to export this information in a way that serves many other possible tasks.

ACKNOWLEDGMENTS

This work is funded by a subcontract with Rensselaer Polytechnic Institute by the Office of Naval Research, under Science and Technology: Apprentice Agents.

REFERENCES

- [1] Keni Bernardin, Koichi Ogawara, Katsushi Ikeuchi, and Ruediger Dillmann. 2005. A sensor fusion approach for recognizing continuous human grasping sequences using hidden Markov models. *IEEE Transactions on Robotics* 21, 1 (2005), 47–57.
- [2] Cynthia Breazeal, Guy Hoffman, and Andrea Lockerd. 2004. Teaching and working with robots as a collaboration. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*. IEEE Computer Society, 1030–1037.
- [3] Sylvain Calinon, Florent D’Halluin, Eric L Sauser, Darwin G Caldwell, and Aude G Billard. 2010. Learning and reproduction of gestures by imitation: An approach based on Hidden Markov Model and Gaussian Mixture Regression. *IEEE Robotics and Automation Magazine* 17, 2 (2010).
- [4] Emily B Fox, Michael C Hughes, Erik B Sudderth, Michael I Jordan, et al. 2014. Joint modeling of multiple time series via the beta process with application to motion capture segmentation. *The Annals of Applied Statistics* 8, 3 (2014), 1281–1313.
- [5] Maria Fox, Malik Ghallab, Guillaume Infantes, and Derek Long. 2006. Robot introspection through learned hidden markov models. *Artificial Intelligence* 170, 2 (2006), 59–113.
- [6] Bradley Hayes and Brian Scassellati. 2016. Autonomously constructing hierarchical task networks for planning and human-robot collaboration. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 5469–5476.
- [7] Guy Hoffman and Cynthia Breazeal. 2007. Cost-based anticipatory action selection for human–robot fluency. *IEEE transactions on robotics* 23, 5 (2007), 952–961.
- [8] Chien-Ming Huang and Bilge Mutlu. 2016. Anticipatory robot control for efficient human-robot collaboration. In *Human-Robot Interaction (HRI), 2016 11th ACM/IEEE International Conference on*. IEEE, 83–90.
- [9] Mitsuo Kawato, Yoji Uno, Michiaki Isobe, and Ryoji Suzuki. 1988. Hierarchical neural network model for voluntary movement with application to robotics. *IEEE Control Systems Magazine* 8, 2 (1988), 8–15.
- [10] Richard Kelley, Alireza Tavakkoli, Christopher King, Monica Nicolescu, Mircea Nicolescu, and George Bebis. 2008. Understanding human intentions via hidden markov models in autonomous mobile robots. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*. ACM, 367–374.
- [11] Jim Mainprice and Dmitry Berenson. 2013. Human-robot collaborative manipulation planning using early prediction of human motion. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 299–306.
- [12] Jim Mainprice, Rafi Hayne, and Dmitry Berenson. 2015. Predicting human reaching motion in collaborative tasks using inverse optimal control and iterative re-planning. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 885–892.
- [13] Olivier Mangin, Alessandro Roncone, and Brian Scassellati. 2017. How to be Helpful? Implementing Supportive Behaviors for Human-Robot Collaboration. (2017). arXiv:1710.11194
- [14] Duy Nguyen-Tuong, Matthias Seeger, and Jan Peters. 2009. Model learning with local gaussian process regression. *Advanced Robotics* 23, 15 (2009), 2015–2034.
- [15] Monica N Nicolescu and Maja J Mataric. 2001. Learning and interacting in human-robot domains. *IEEE Transactions on Systems, man, and Cybernetics-part A: Systems and Humans* 31, 5 (2001), 419–430.
- [16] Stefanos Nikolaidis, Przemyslaw Lasota, Gregory Rossano, Carlos Martinez, Thomas Fuhlbrigg, and Julie Shah. 2013. Human-robot collaboration in manufacturing: Quantitative evaluation of predictable, convergent joint action. In *Robotics (ISR), 2013 44th international symposium on*. IEEE, 1–6.
- [17] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. 2015. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 189–196.
- [18] Koichi Ogawara, Jun Takamatsu, Hiroshi Kimura, and Katsushi Ikeuchi. 2002. Modeling manipulation interactions by hidden Markov models. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, Vol. 2. IEEE, 1096–1101.
- [19] Lawrence Rabiner and B Juang. 1986. An introduction to hidden Markov models. *iee assp magazine* 3, 1 (1986), 4–16.
- [20] Aleksandar Vakanski, Iraj Mantegh, Andrew Irish, and Farrokh Janabi-Sharifi. 2012. Trajectory learning for robot programming by demonstration using hidden Markov model and dynamic time warping. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42, 4 (2012), 1039–1052.
- [21] S. Zeylikman, S. Widder, A. Roncone, O. Mangin, and B. Scassellati. 2017. The HRC model set for human-robot collaboration research. (2017). arXiv:1710.11211