Adaptive Planning with Evidence Based Prediction for Improved Fluency in Routine Human-Robot Collaborative Tasks

Christopher K. Fourie
Computer Science and Artificial Intelligence Laboratory (CSAIL)
Massachusetts Institute of Technology
77 Massachusetts Ave, Cambridge, MA, 02139

Introduction

Robotics has seen widespread use in manufacturing environments, particularly in the automotive sector, where their introduction has alleviated costs, reduced workload and improved throughput. However, limitations exist in what can be automated, and tasks remain in which the human is not only required, but of crucial importance to the success of the task. In such scenarios, a dynamic interface exists between humans and machines, in which a form of collaboration is required to ensure effective and fluent interaction. From the perspective of the human, the effectiveness of this collaboration can be assessed on several criteria, such as trust and fluency (Hoffman 2013).

Many of the complex scenarios in this domain largely fall under the concept of a joint action workspace (Sebanz, Bekkering, and Knoblich 2006), where a human and a robot are acting together to accomplish a common goal, in a manner that requires them to interact with and influence one another. These principles have been the subject of ongoing investigation - for example, in (Nikolaidis et al. 2017), a game-theoretic approach is taken in which the robot approximates the human's perception of its capabilities, while in (Dragan 2017), a series of mathematical and probabilistic models is discussed for describing human states and actions, often in response to robot action.

In short, an interpretation of the consequences of a particular action must exist, which either partner must extrapolate upon or use to adapt to the other.

This is particularly true in repetitive tasks as these help ensure that a learned model of such behavior remains stationary and interactions remain fluent: in (Harrison et al. 2003), it was found that being familiar with a teammate facilitated performance, and, more importantly, that patterns in task execution persisted in future iterations, despite changes in deadline or interventions. This would suggest that performance in repetitive tasks converges to a steady state behavior that is robust to variations.

This thesis work intends to explore the development of a shared mental model between an autonomous agent and a human, where we aim to promote fluency in continuing interactions defined by repetitive tasks. That is, with repetitive actions, experimentation and increasing iterations, we wish the robot to learn how its own behavior affects that of its partner. To accomplish this, we propose a model that encodes both human and robot actions in a probabilistic space describing the temporal transition points between activities. The purpose of such a model is not only in passive predictive power (understanding the future actions of an associate), but also to encode the *latent effect of a robot's action on the future actions of the associate*.

Further, such a model raises some interesting possibilities: firstly, the development of a planning algorithm that takes advantage of such a model could, in principle, be used to improve fluency in repetitive tasks. Secondly, we believe that the combination of such an algorithm and model could be used to help a team in training meet long-term behavior goals and requirements. In this thesis, we wish to focus on the specific case of a collaborative robotic partner in a joint-action framework, in which the robot is an active participant in the task and working toward the same goal as the human.

Thesis Plan

Recent Work: Practical Demonstration of an Assistive Robot on a Factory Floor

Previous work with collaborators (Iqbal et al. 2018) focused on the development of a robotic assistant on an automotive factory floor, tasked with bringing assembly parts to an associate for installation. The project was a demonstration of specific capabilities: primarily, it was a showcase of a state-of-the-art activity recognition system (Hayes and Shah 2017), but it also included practical additions required for an actual demonstration. These additions included the implementation of the required robotic behavior (collecting and handing over the part, the work of my colleague), and a dynamic scheduling module to ensure that the part was handed over at the correct time (my personal contribution to the project).

For the project, we needed to develop a temporal model from training segments x_j and labels $l_j = f(x_j)$ that predicted the times t_j that handovers would need to take place, assuming that activity ordering was known. We needed to create an appropriate schedule to ensure handovers took place at the specified times.

I used a Gaussian Mixture Model for the temporal model in which the task durations for each activity, $\Delta t_i, i \in \{1, 8\}$ were modeled as the feature vector. Once the model was trained, I analytically computed the conditional distribution of the model with the observed evidence and utilized the largest peak of the resulting joint distribution as the most likely setting of the future activity durations. I then used knowledge of the (fixed) activity ordering to construct future time point predictions. The schedule became linear and the latest time of robot action initiation could be directly inferred in polynomial time.

This approach was noted to have several advantages: it was fast to train, inference was computationally cheap, and it was scalable. The model was also capable of capturing multiple modes and variations in associate behavior and could easily adapt to subtle variations in performance.

However in the context of repetitive tasks, the approach had some limitations. In particular, the activities were constrained to a particular ordering, but not as a requirement of the task - this was largely a requirement of the activity recognition, in which activity segmentation became intractable with multiple hypotheses. However, with the knowledge that repetitive tasks would tend to result in a consistent ordering, it would be desirable to learn the ordering online per a particular associate, and utilize this information to disambiguate low-level activity recognition.

Current Work: Integrated Activity Recognition and Temporal Activity Prediction

Given a set of training examples, composed of a trajectory x_i and ground truth activity labels $l_i = f(x_i)$, we seek to develop a system that predicts the times associated with the start and end points of the activity segments, $[t^s_{ai}, t^e_{ai}]$. In particular, such a model should be able to directly encode ordering.

We are currently designing an integrated system that combines a temporal prediction model with activity recognition. The temporal prediction model is based on the GMM-based temporal predictor discussed in the previous section while the activity recognition is based off the work in (Hayes and Shah 2017). We hypothesize that these two elements can be mutually beneficial, as the temporal prediction can be used to disambiguate multiple hypotheses in activity recognition while short-term goal prediction could be used to update temporal predictions.

Pending the outcome of remaining development and extensive testing, we hope to submit this work for publication in December, 2018.

Future Work: Integrated Prediction and Planning for Improved Fluency in Human-Robot Interaction

Through integrating the time of robot action (t^r_i) into the temporal predictive model, we seek to use the model to optimize robot actions to minimize delays and promote fluency.

We currently intend to formulate the problem as a constraint optimization, but believe that a modeling relaxation may exist to improve tractability and efficiency of computation.

Future Work: Human Experiments and an Investigation into the Controlled Evolution of Fluency and Efficiency in Repetitive Tasks

An autonomous agent could utilize its predictive model to influence both human and robot action and guide the team towards system goals. Such behavior should be integrated in training so that persistent long-term patterns meet system goals.

Human teams evolve to be more efficient and fluent in multiple iterations of the same task, but the work of (Harrison et al. 2003) would suggest that patterns persist post-training. An autonomous agent may be able to influence training in such a way as to steer persistent behavior towards the joint action goals. The implication of this is that such a system could potentially help associates meet their long-term targets (eg. efficiency or time constraints) while maintaining fluent interactions in training.

Conclusion and Timeline

I would like to orientate my doctoral research around human-robot teams in manufacturing environments and around the idea that repetitive task environments are a key domain to be focusing on. I plan to complete the first contribution, the current work, by the time of the doctoral consortium, and the future contributions within 6-8 month blocks succeeding it (aiming for an expected graduation date of 2021).

References

Dragan, A. D. 2017. Robot planning with mathematical models of human state and action. *arXiv preprint arXiv:1705.04226*.

Harrison, D. A.; Mohammed, S.; McGrath, J. E.; Florey, A. T.; and Vanderstoep, S. W. 2003. Time matters in team performance: Effects of member familiarity, entrainment, and task discontinuity on speed and quality. *Personnel Psychology* 56(3):633–669.

Hayes, B., and Shah, J. A. 2017. Interpretable models for fast activity recognition and anomaly explanation during collaborative robotics tasks. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, 6586–6593. IEEE.

Hoffman, G. 2013. Evaluating fluency in human-robot collaboration. In *International conference on human-robot interaction (HRI), workshop on human robot collaboration*, volume 381, 1–8.

Iqbal, T.; Li, S.; Fourie, C. K.; Hayes, B.; and Shah, J. A. 2018. Fast online segmentation of activities from partial trajectories. Submitted for publication in RA-L.

Nikolaidis, S.; Nath, S.; Procaccia, A. D.; and Srinivasa, S. 2017. Game-theoretic modeling of human adaptation in human-robot collaboration. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '17, 323–331. New York, NY, USA: ACM.

Sebanz, N.; Bekkering, H.; and Knoblich, G. 2006. Joint action: bodies and minds moving together. *Trends in cognitive sciences* 10(2):70–76.