

Integrating Planning and Recognition to Close the Interaction Loop

Richard G. Freedman

College of Information and Computer Sciences
 University of Massachusetts Amherst
 freedman@cs.umass.edu

Introduction

In many real-world domains, the presence of machines is becoming more ubiquitous to the point that they are usually more than simple automation tools. As part of the environment amongst human users, it is necessary for these computers and robots to be able to interact with them reasonably by either working independently around them or participating in a task, especially one with which a person needs help.

This interactive procedure requires several steps: recognizing the user and environment from sensor data, interpreting the user's activity and motives, determining a responsive behavior, performing the behavior, and then recognizing everything again to confirm the behavior choice and re-plan if necessary. At the moment, the research areas addressing these steps, activity recognition, plan recognition, intent recognition, and planning, have all been primarily studied independently. However, pipelining each independent process can be risky in real-time situations where there may be enough time to only run a few steps. This leads to a critical question: *how do we perform everything under time constraints?* In this thesis summary, I propose a framework that *integrates these processes* by taking advantage of features shared between them. This includes my current work towards this preliminary system and outlines how I plan to complete the integration for time-constrained interaction.

Background

One of the earliest areas of artificial intelligence, *planning* is the study of automated action selection. Early approaches usually involved representing the world as a list of logic statements and searching for a sequence of actions which would modify the list until it contained the set of goal conditions; the notation used for this is called *STRIPS*. Modern approaches range from improving search over STRIPS to decision theoretic planning with MDPs and its variants to approximation methods to handle uncertainties in the world.

As its inverse problem, *plan recognition* (PR) tries to identify the problem an agent is solving given its observed actions. The actions and problems are usually represented at a higher level such as STRIPS. *Activity recognition* (AR) works at the lower level by interpreting sensor data as higher-level actions. *Intent recognition* (IR) tries to predict the agent's specific goal or upcoming actions which allows

some degree of foresight into the observed agent's behavior. Collectively, these fields of recognition are referred to as *PAIR* and have become a more popular area of research recently, including the topic of a Dagstuhl Seminar (Goldman et al. 2011). Various problems in PAIR are studied in other fields which has made the literature vast, but they are still studied independently or pipelined in most these works.

Simultaneous Plan and Activity Recognition

The formulation of a typical recognition problem is as follows: given a sequence \mathcal{O} of observations o_1, o_2, \dots, o_n , determine which task(s) in library L the agent is performing. In AR, each o_i is a sensor reading and L is a set of actions or activities. Supervised machine learning and graphical models are usually used to infer the label in L which best describes \mathcal{O} . For PR, each o_i is a STRIPS action and L contains the set of all tasks. A weighted matching method is often used to compare \mathcal{O} to some of each task's solutions called *plans*, a sequence of STRIPS actions satisfying the goal conditions. A recent method used to perform this matching is parsing hierarchical task networks, revealing a *parallel between PR and natural language processing* (NLP) (Geib and Steedman 2007). \mathcal{O} is like a sentence and the breakdown of a complex task into subtasks is like a grammar of derivations.

I began the extension of this analogy by considering another problem in NLP: topic modeling. Unlike parsing which determines the underlying structure of a sentence, topic modeling investigates the concepts discussed in a collection of documents by finding clusters of related words called *topics*. For the popular Latent Dirichlet Allocation (LDA) topic model (Blei, Ng, and Jordan 2003), each topic is simply a distribution over the set of words and each document is a distribution over the set of topics; the distributions are learned using unsupervised learning to model the training data. For interaction in a variety of domains, an unsupervised approach is more appealing because it is possible to learn activities in each domain as it is added to the system. Furthermore, LDA's *bag-of-words assumption* where each o_i is independent of the rest of \mathcal{O} was appealing to begin integrating AR and PR due to the mismatch of the two sequence formations. A sensor records over time so that a single action has multiple consecutive o_i for AR, but a single STRIPS action is only one o_i for PR. Figure 1 illustrates how this analogously treats actions like topics. Each word is a sensor reading and the distribution over these readings




Figure 1: Analogy of Distributions for Topics and Actions

describes an action while the task of the recording session is represented by the distribution of actions. *Hence inferring a topic with LDA performs AR and the distribution of the inferences enables us to approximate PR simultaneously.*

After determining how to map RGB-D sensor readings of three-dimensional stick figures to a space of words, I ran LDA on a small dataset. The results provided evidence supporting my hypothesis since each topic contained postures resembling simple actions as in Figure 2 (Freedman, Jung, and Zilberstein 2014). We then extended LDA to consider the presence of nearby objects and/or global temporal patterns and found evidence that the information provided by these two factors are not only independent, but assist disambiguating actions which contain common postures (Freedman, Jung, and Zilberstein 2015). In future work, we will investigate additional variations and their insights for PAIR.

Integration of Planning with Plan Recognition

Ramírez and Geffner (2010) introduced a compilation of PR problems into classical planning problems. It assigned a distribution over L where Bayes's Rule weighs the plans against each other using the most optimal plans with and without \mathcal{O} as a subsequence. This accounts for the probability of the agent solving each task conditioned on its observed actions, considering optimal (shorter) plans to be more likely. While the accuracy for the method is very strong, a temporal plot of the probabilities showed that it only achieved this accuracy towards the *completion of the plan* when the final actions were observed.

While their compilation is excellent when all the observed actions are available, this is not as practical in interactions because the observed agent will likely be in mid-execution of a plan. In collaboration with Fukunaga (U. Tokyo), I have proposed two approaches to address this (Freedman and Fukunaga 2015). The first one generates a *dynamic prior* for Bayes's Rule that removes the bias for shorter plans and converges to the true prior over time. The second one counts the number of permutations of a partially-ordered plan in order to account for the number of optimal plans rather than their length alone. The updated distribution over L can be used to aggregate the lists of logic statements which must be true for each goal and identify the most necessary conditions. If a set of conditions is shared between the most likely tasks, then satisfying them should be required regardless of the task. Thus a second pass of the planner should yield a plan which the machine may execute to interact properly even if the recognized task is still ambiguous. I developed a simulator for a multiplayer version of the Sokoban game (a benchmark in planning research) and we will begin implementing and testing these methods within the next two months.




Figure 2: Learned Clusters of Postures Resembling Sitting

Status of Thesis

I plan to close the interactive loop by completing the integration of these processes. Besides continuing the works above, there are several key remaining tasks. The most important one is bridging the gap between simultaneous PR and AR and the integration of PR with planning. Although it seems trivial because both contain PR, they do not align due to the unsupervised nature of topic models. Recognized actions are clusters of postures without annotation while actions in the newer research are assumed to be in STRIPS which is designed by humans. I have begun to identify methods for autonomously extracting features of the postures with the greatest probability mass in each cluster and using them to describe the respective action (Freedman and Zilberstein 2015). With Wallach, I am exploring analogies in topic modeling. I have also spoken with Fukunaga about using constraint optimization to align LDA clusters with STRIPS operators using ordering of each \mathcal{O} and these extracted features. After implementing and testing, it will be ideal to integrate IR to better predict upcoming actions so that the machine does not interfere with the observed user. For this, I intend to investigate the planning graph and determine how to probabilistically select action nodes which are more likely to be executed. Once these are in tact, the preliminary interaction loop will be complete and optimization will be necessary to make it usable under realistic time constraints.

References

- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent Dirichlet allocation. *Jour. of Machine Learning Research* 3:993–1022.
- Freedman, R. G., and Fukunaga, A. 2015. Integration of planning with plan recognition using classical planners (extended abstract). In *AI for HRI: Papers from the AAI 2015 Fall Symp.*, 51–59.
- Freedman, R. G., and Zilberstein, S. 2015. Automated interpretations of unsupervised learning-derived clusters for activity recognition. In *Wkshp. on Learning for Human-Robot Collaboration*.
- Freedman, R. G.; Jung, H.-T.; and Zilberstein, S. 2014. Plan and activity recognition from a topic modeling perspective. In *Proc. of the 24th Int'l Conf. on Auto. Plan. and Sched.*, 360–364.
- Freedman, R. G.; Jung, H.-T.; and Zilberstein, S. 2015. Temporal and object relations in unsupervised plan and activity recognition. In *AI for HRI: Papers from the AAI 2015 Fall Symp.*, 48–50.
- Geib, C. W., and Steedman, M. 2007. On natural language processing and plan recognition. In *Proc. of the 20th Int'l Joint Conf. on AI*, 1612–1617.
- Goldman, R. P.; Geib, C. W.; Kautz, H.; and Asfour, T. 2011. Plan Recognition – Dagstuhl Sem. 11141. *Dagstuhl Reports* 1(4):1–22.
- Ramírez, M., and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *Proc. of the 24th AAI Conf. on AI*.