

A Unifying Perspective of Plan, Activity, and Intent Recognition

Richard G. Freedman

Smart Information Flow Technologies, LLC

rffreedman@sift.net

Shlomo Zilberstein

University of Massachusetts Amherst
College of Information and Computer Sciences
{freedman, shlomo}@cs.umass.edu

Abstract

Plan, activity, and intent recognition (PAIR) have much in common and are often discussed within the same research community. However, the actual methods for each type of recognition are rarely considered with respect to the others. This leads to lost opportunities where the work in one type of PAIR can play a role in the advancement of other types. To motivate taking advantage of commonalities between different types of PAIR, we propose a generalized formalization of recognition that can be used in all three types of recognition. We also propose a perspective that emphasizes how PAIR applies to different aspects of the generative planning and execution process. The relations between PAIR in this perspective enable us to illustrate a few examples of how current PAIR methods can be used in other types of recognition outside the one for which they were developed.

1 Introduction

Several decades since the establishment of automated planning as an area of study, the artificial intelligence (AI) community has branched into various specialized communities that pursue specific challenges that have been identified within the scope of AI. Quite a few of them began to focus on perception tasks, and the majority of the early research on perception was strictly interpreting sensor data. With much inspiration from the application of intelligent robotics, this led to a lot of research within areas such as computer vision (Thorpe et al. 1988) and tactile sensing (Harmon 1982). However, perception involving higher-level thought was eventually realized in story analysis (Kautz 1991) and understanding user activities within program applications (Madani, Bui, and Yeh 2009).

The most distinguishing feature between these two forms of perception is the *observation focus*. The former, lower-level perception aims to obtain an overall environment description; this creates a world state in which the perceiving agent can act. The latter, higher-level perception is associated with understanding agents and their underlying decision making processes. The focus on observing other agents independently of the environment itself adds a “black box” around their actions that must be interpreted in addition to rules of the environment (such as physics), which all affect

how the perceiving agent can act in the world. The formal establishment of the plan, activity, and intent recognition (PAIR) community studying this happened within the past decade (Goldman et al. 2011) following a short series of workshops originally titled “Modeling Other Agents from Observations” (MOO). The workshop series was eventually renamed PAIR due to the common methods and themes in presented research, but there are still discrepancies regarding what PAIR specifically studies in each type of recognition.

One goal of this manuscript is to begin addressing these discrepancies and *investigate how all the types of recognition are interrelated*. Despite their common origin and being discussed at the same venues, they are not often discussed with respect to the other types of recognition (though a recent compendium does mention in the preface that plan and intent recognition are strongly related (Sukthankar et al. 2014)). Each form of recognition is considered to be different to some degree, but these degrees have not been formally described outside qualitative descriptions. Section 2 will introduce our proposed formalization of recognition as a whole and respectively describe each type of recognition in PAIR with it. This will ideally establish a basis for comparing and contrasting each type of recognition, and we invite the PAIR community to revise and extend it.

The second goal of this manuscript is to provide a context that illustrates how each aspect of PAIR is related to the overall task of perceiving other agents. This perspective can serve as a means for the formalization above to better identify when techniques developed for one type of recognition can be used to aid or perform other types as well. We describe this perspective of PAIR in Section 3, which is based on both the generative decision making process that derives a plan or policy for how to act and the execution process that performs the selected actions. Section 4 follows with a few examples that we identified where past PAIR research can be applied to other types of recognition. We also invite the PAIR community to identify additional examples and further investigate how their work can play a role in other types of recognition besides the one of their focus.

2 Background and Formalization

As Section 1 briefly mentions, there are relations between the AI specializations of PAIR and automated planning that we will explore in this manuscript. We thus provide a brief

description on planning and each type of recognition, using our proposed general definition:

Definition 1 A recognition problem is a tuple $\mathcal{R} = (\mathcal{KB}, \mathcal{O}, \mathcal{H})$ where \mathcal{KB} is a knowledge base representing what the observing agent knows (and thus what information is at its disposal), \mathcal{O} provides a sequence of observations, and \mathcal{H} lists the possible hypotheses that can be recognized.

For each type of recognition, we will describe the typical values that it uses for the elements of \mathcal{R} . This will serve as a common ground for PAIR’s formalizations, which broadly resemble matching \mathcal{O} to some element of \mathcal{H} based on \mathcal{KB} . \mathcal{KB} and \mathcal{H} are often intertwined such that the question “is $h \in \mathcal{H}$ the correct match?” can be answered using a portion of \mathcal{KB} that relates to just h . \mathcal{KB} is specifically defined with respect to the observing agent that performs recognition, R_{ing} . Any information regarding the observed agent that is recognized, R_{ed} , is only in \mathcal{KB} if R_{ing} knows or assumes it.

Automated Planning

One of the earliest challenges posed to the AI community involved machines being capable of making decisions autonomously at or above the level of human experts. This led to the establishment of the automated planning and scheduling community that particularly studies representation of tasks, problem solving under various conditions ranging from uncertainty to resource constraints, and higher-level decision making processes such as metareasoning.

Definition 2 A planning problem is a tuple $\mathcal{P} = (\mathcal{D}, I, G)$ where \mathcal{D} is a domain that models the world, I provides the initial setup of the world, and G lists the task’s completion conditions.

Definition 3 A planning domain models the world in which the agents act. The contents of its tuple should describe the set of states S and the set of actions A that transition between states.

There are many different automated planning formalisms that each represent some problem aspects more efficiently than others, which allows there to be different classes of problems with specialized algorithms that can effectively solve them. Common representations include various orders of logic (Gerevini et al. 2009; Bäckström 1992; Pednault 1989), task hierarchies (Erol, Hendler, and Nau 1994), and degrees of uncertainty (Bellman 2003; Younes and Littman 2004; Sanner 2010; Fu et al. 2016). Depending on the amount of certainty in the problem’s formulation, there are several types of solutions that determine how the agent(s) should act to solve the problem.

Definition 4 A plan π is a sequence of actions $a_1, a_2, \dots, a_{|\pi|} \in A$ such that each action a_i is applicable in the state derived from performing the previous actions in order from the initial state $a_{i-1}(a_{i-2}(\dots a_1(I) \dots))$ and the resulting state after completing all the actions in order satisfies the goal conditions. If the actions do not have costs, then the cost of π is $cost_\pi = |\pi|$. If the actions do have costs, then the cost of π is the total action cost $cost_\pi = \sum_{i=1}^{|\pi|} cost_{a_i}$.

Definition 5 A policy is a function $\pi : S \rightarrow 2^A$ that instructs what action to perform at each state in the state space. The policy is deterministic iff $|\pi(s)| \leq 1$ for all $s \in S$ and non-deterministic otherwise (Cimatti et al. 2003).

The majority of automated planning solvers use a search technique to find a plan or policy for \mathcal{P} . The specific space that is searched and how it plays a role in finding the solution varies per algorithm.

Plan Recognition

Plan recognition is viewed as an inverse problem to automated planning because the observation sequence is composed of executable actions. That is, each $o \in \mathcal{O}$ often satisfies $o \in A$. However, \mathcal{O} is not often complete; R_{ing} might ‘blink’, a certain action might not be observable with the available sensors, or a transcription might only point out ‘key actions’ rather than provide a ‘play-by-play’ summary. Thus the primary challenge in plan recognition is to answer the question, “what is the agent doing overall?” \mathcal{H} is usually a set of plans or high-level tasks. However, the answer to the question is paired with the secondary challenge that has received more emphasis lately (Geib 2009; Mirsky, Gal, and Shieber 2017): “why is this the correct answer?”

Definition 6 An explanation is a plan or policy π that best resembles \mathcal{O} ’s sequence and justifies why some $h \in \mathcal{H}$ is the answer to plan recognition problem \mathcal{R} .

Depending on \mathcal{KB} , the explanation may have contingencies, be a task network, or more. Overall, the structure of, and information in, the knowledge base has the most impact on the different plan recognition algorithms. Two common knowledge bases are based on how much R_{ing} knows about R_{ed} and the world in which they are acting.

Definition 7 A plan library is a type of knowledge base that contains precomputed plans for solving some set of automated planning problems. This includes a grammar that can construct plans without solving the automated planning problems.

Definition 8 A planning domain is a type of knowledge base that models the world in which the agents act. This is identical to Definition 3, but is now a tool for “thinking in another agent’s shoes” rather than personal decision making.

The knowledge base was partitioned with plans for each possible hypothesis in earlier research, which became plan libraries (Kautz 1987). While libraries are still used and have their advantages, knowledge bases are more expressive and generalized when the domain itself is provided (Ramírez and Geffner 2009). In particular, \mathcal{H} can change independently of updating \mathcal{KB} as long as the new hypotheses are solvable in the original domain. Unlike automated planning, it is much harder to define the case where there is no solution in \mathcal{H} . Some algorithms return a distribution over \mathcal{H} rather than returning a subset of elements in \mathcal{H} (Ramírez and Geffner 2010), but this is still a relative comparison that implies more ambiguity as the distribution over some subset of \mathcal{H} becomes uniform. The spread

metric $S_{\%}$ (Ramírez and Geffner 2010) and quality metric $Q_{\%}$ (E-Martín, R-Moreno, and Smith 2015) have been proposed to quantify the number of most-likely hypotheses in $\mathcal{H}_{\%} = \{h \in \mathcal{H} \mid P(h|\mathcal{O}) \text{ is at the } \% \text{ percentile or greater}\}$ and how often $\mathcal{H}_{\%}$ contains the correct answer respectively. Ideally, for greater values of $\%$, $Q_{\%}$ will be large and $S_{\%}$ will be small.

Methods used to solve plan recognition problems vary widely. The early works simply matched sequences based on the available plan library (Kautz 1987), and they were extended through tie-breaking methods like abduction (Charniak and Goldman 1991) and weighted clauses (Hobbs et al. 1993; Inoue and Inui 2011). Task hierarchies later became the automated planning representation of choice (Geib and Steedman 2007), which yielded parsing approaches (Geib and Goldman 2009; Mirsky, Gal, and Shieber 2017). Recognition as planning uses logic-based planning representations and runs an automated planning solver to simulate plans that comply with \mathcal{O} (Sohrabi, Riabov, and Udrea 2016).

Activity Recognition

Activity Recognition is the process of identifying a single action or task given a sequence of lower-level observations. Lower-level is relative to \mathcal{H} 's contents. Task hierarchies and logic-based representations with complex tasks and/or macroactions in \mathcal{H} often observe the simpler tasks that make up larger ones so that each $o \in \mathcal{O}$ again satisfies $o \in A$. However, it is also possible that $\mathcal{H} \subseteq A$ so that the high-level information is a single action; then the low-level information in \mathcal{O} might range from sensor data to environment configurations. In both cases, the challenge in activity recognition is to answer the question, “what is the agent doing *at the moment*?”

The majority of the solutions to activity recognition involve either probabilistic models or machine learning classifiers (Anjum and Ilyas 2013; Gori et al. 2015). In particular, they commonly use or extend traditional time-series models that rely on latent state information that depends on up to one previous state.

Definition 9 A hidden Markov model (HMM) is a probabilistic model with random variables $X_0, X_1, \dots, X_n \in V$ and $O_1, O_2, \dots, O_n \in \Omega$ where X_i form a Markov chain and each O_i 's value only depends on the value of X_i . The O_i random variables are all observed, but the X_i random variables are not observed.

Definition 10 A bag of words (BOW) model is a probabilistic model with random variables $X_0, X_1, \dots, X_n \in V$ and $O_1, O_2, \dots, O_n \in \Omega$ where all X_i are sampled from some global distribution independently of each other and each O_i 's value only depends on the value of X_i . The O_i random variables are all observed, but the X_i random variables are not observed.

When \mathcal{H} 's elements use higher-level action representation, the model is often extended to account for its composition from smaller tasks (Fine, Singer, and Tishby 1998; Bui, Venkatesh, and West 2002; Bui, Phung, and Venkatesh 2004). On the other hand, \mathcal{H} that contain simpler action representations just use out-of-the-box HMM and BOW models

on the raw sensor input (Jung et al. 2015; Chikhaoui, Wang, and Pigot 2012; Rieping, Englebienne, and Kröse 2014; Chen, Diethel, and Flach 2016). However, some research has looked at techniques for processing sensor data into alternative representations that either take advantage of domain information or embed more complex temporal-spatial features; then the processed observations are used in a standard classifier or model (Huỳnh, Fritz, and Schiele 2008; Zhang and Parker 2011; Freedman, Jung, and Zilberstein 2014).

For activity recognition, \mathcal{KB} usually stores the information needed to generate and run the classifier or probabilistic model. This includes the probability tables (defined by Bayesian priors or computed with training data) and possible values for each variable, including the observed and hidden ones.

Intent Recognition

Intent Recognition studies the ‘driving motivation’ behind what is observed. In most cases, the underlying motivation for doing something has been characterized as a subset of possible goal conditions; logic-based representations define $\mathcal{H} \subseteq 2^F$ where F is a set of features that can describe states in S and task hierarchies define \mathcal{H} as a set of high-level tasks (much like plan recognition with task hierarchies). This is the reason that the term goal recognition is also used sometimes. However, goals themselves can be selected for various reasons. We are not aware of any research on recognition at the metareasoning level, but Callaway et al. (2017) trace users’ decision making processes to study their planning and information-gathering strategies. Without approaching metareasoning for decision making, it is viable that some actions are performed for the sake of facilitating another action that is crucial to R_{ed} 's plan. This resembles taking actions to repair missing causal links that are required to perform upcoming actions in a plan (Levine and Williams 2014). If these reparation actions are not useful for the goal’s completion, though, then there might be a misinterpreted motive. In this case, intent recognition can be viewed as a prediction problem instead where $\mathcal{H} \subseteq A$ or $\mathcal{H} \subseteq S$ (Baker, Saxe, and Tenenbaum 2009).

Though there is potential overlap with plan recognition’s formulation, the fundamental question is different: “why is the agent doing this?” The variation of recognition as planning cited in the Plan Recognition subsection above modified \mathcal{H} from sets of goal conditions to sets of plans, but others made variations for intent recognition with a focus on only identifying the goal conditions rather than the plan or policy explanations (Ramírez and Geffner 2011; E-Martín and Smith 2017; Pereira, Oren, and Meneguzzi 2017). In all the intent recognition approaches mentioned so far, each observation $o \in \mathcal{O}$ is again an executable action such that $o \in A$. However, prediction based on agents’ trajectories in continuous space (Unhelkar et al. 2015; Vered and Kaminka 2017) is usually observed as either spatial coordinates such as $o \in \mathbb{R}^2$ or kinematic configuration spaces consisting of rotations and translations of joints (Mainprice, Hayne, and Berenson 2016).

As with plan recognition and activity recognition, \mathcal{KB} will

contain the underlying information to construct any models that are used in the intent recognition algorithm. Though the specific implementations are different, they are typically probabilistic models, plan libraries, and/or plan domains. Furthermore, a newer area of research inspired by intent recognition problems is goal recognition design (Keren, Gal, and Karpas 2014; 2015), which studies how to optimize the performance of intent recognition algorithms by modifying the environment (and thus \mathcal{KB}). The approaches prune grounded actions from A that are not mandatory to reach a goal in \mathcal{H} and are ambiguous when R_{ing} observes R_{ed} performing them.

Definition 11 Let \mathcal{D} be an automated planning domain, I be an initial state, and \mathcal{G} be a set of goal conditions; then we can define automated planning problems $\mathcal{P}_i = (\mathcal{D}, I, G_i \in \mathcal{G})$ with their respective set of optimal plan solutions Π_i^* . A sequence of actions π is non-distinctive if there exists unique problems \mathcal{P}_1 and \mathcal{P}_2 (that is, $G_1 \neq G_2$) such that π is a prefix of both some optimal plan $\pi_1^* \in \Pi_1^*$ and some other optimal plan $\pi_2^* \in \Pi_2^*$.

Definition 12 The worst case distinctiveness (WCD) is the greatest amount of ambiguity possible between optimal plans for two different goals. That is, $wcd(\mathcal{D}, I, \mathcal{G}) = \max_{\pi \in \Pi_{\mathcal{D}}^{I \rightarrow \mathcal{G}}} |\pi|$ where $\Pi_{\mathcal{D}}^{I \rightarrow \mathcal{G}}$ is the set of all non-distinctive paths for \mathcal{D} , I , and \mathcal{G} .

Goal recognition design with respect to stochastic action outcomes (Wayllace et al. 2016) uses similar definitions to those above with policies instead of plans.

3 A Unified Perspective of Recognition

In order to focus on what each type of recognition shares and how they complement each other, we need a single context in which they all apply. As Definition 1 and Section 1’s general history of PAIR both imply, PAIR revolves around understanding other agents from their actions. Thus we will consider the generative planning context in which R_{ed} selects actions that are then performed while R_{ing} observes:

1. We assume that R_{ed} generally uses some automated planning algorithm \mathcal{A} to solve its problems. We will also assume that R_{ed} has some degree of awareness so that R_{ed} knows about the surrounding environment and can reasonably identify the current state of the world.
2. Suppose that R_{ed} now has a goal G that must be satisfied, either directly assigned by some agent or derived from some personal desire (Riedl and Young 2010) or interactive task (Freedman and Zilberstein 2017).
3. Using the assumptions, R_{ed} is able to construct an automated planning problem \mathcal{P} using the current world state (or belief if there is partial observability) for I , knowledge about the environment for \mathcal{D} , and the given G .
4. R_{ed} now uses \mathcal{A} to solve \mathcal{P} . Depending on \mathcal{A} , the plan or policy π can be a sequence of actions, have contingencies, only map a subset of the state space, be a task network, etc.; what matters most is that R_{ed} now has instructions for how to act.

This process provides insight into two components of recognition problems: \mathcal{KB} and \mathcal{H} . Specifically, the *knowledge base is primarily built upon the assumptions about R_{ed} and the hypotheses are derived from the possible ways that a goal can be assigned*. Because we do not know exactly what R_{ed} knows or is thinking (if we did, then recognition would be a more trivial task), these are designed for R_{ing} to use as models for each observed agent. Interactive systems that combine recognition and planning sometimes impose R_{ing} ’s own knowledge on those it observes (Levine and Williams 2014; Freedman and Zilberstein 2017) such as shared planning algorithms and/or plans, and others use unique models for R_{ing} and R_{ed} (Geib et al. 2016). However, the thought process alone is not often sufficient for the recognition problem because R_{ing} needs to observe the actions itself. This progresses the process above to R_{ed} ’s execution of actions and how they affect the world state:

5. Without any form of intervention, only the ‘natural physics’ of the environment apply as a *closed system*. We can represent this closed system as a Markov chain because it is usually not the case that every physics calculation is accounted for. In a deterministic world where they are perfectly predictable, each row of the transition matrix is simply a vector of all 0’s except for a single entry with value 1 for the guaranteed state transition.
6. When R_{ed} performs a particular action $a \in A$ that manipulates the environment, a ’s effects will change the world state with respect to its own transition matrix. As a generalized case, consider a corresponding Markov decision process’s state transition function T that yields the distribution over states given a current state and the action performed. Then we have a new Markov chain for the environment while this action is being performed; this should also account for the intervention of ‘natural physics’ alongside R_{ed} ’s action. If the action is deterministic alongside the world, then the transition matrix rows will again be vectors of all 0’s and a single 1.
7. Observations are restricted by R_{ing} ’s available sensors. They might provide information about some part of the current state such a subset of state features F (Sohrabi, Riabov, and Udrea 2016) (hardware sensor readings or environment descriptions), the action that was performed (Kautz 1991) (a story or broadcast), or R_{ed} ’s underlying thought process (Mirsky and Gal 2016; Geib et al. 2016; Fox, Long, and Magazzeni 2017) (querying the agent directly).

The observations about the other agents and/or the environment are thus dependent on R_{ing} ’s sensing capabilities and the state of the world. This implies that the sequence of observations \mathcal{O} are constrained snapshots of the world state. Because the state transitions with respect to one of many Markov chains, we can view a short-term version of this context as a HMM. R_{ed} ’s current action $a_r \in \pi$ determines which transition matrix to apply to the latent states, and R_{ing} extracts some information from these states to get some $o_i, \dots, o_{i+j} \in \mathcal{O}$. Then the action changes to a_{r+1} based on what π says to do next, which also changes the transition matrix between the latent states while R_{ing} per-

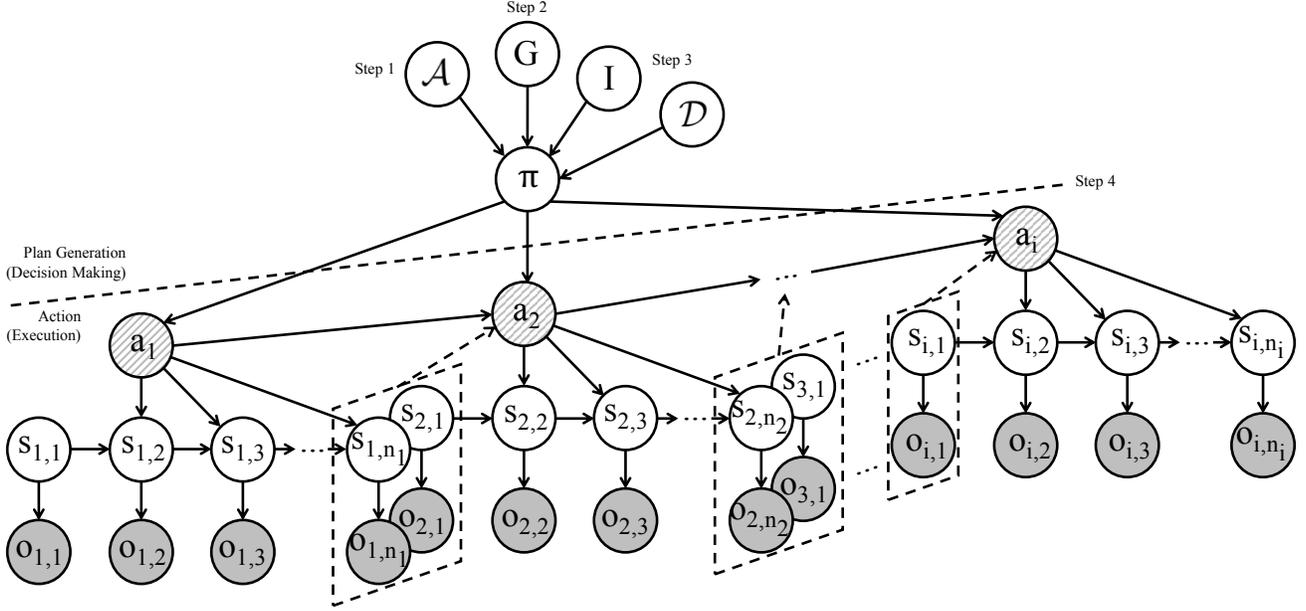


Figure 1: A visualization of R_{ed} 's decision making and action execution processes. Directed edges indicate a dependency between two components, but dependencies are not necessarily probabilistic. The dashed boxes indicate the same variable under two different names to bridge the HMMs: $s_{k,n_k} = s_{k+1,1}$. Shaded nodes indicate observations. The partially shaded action nodes a_1, a_2, \dots, a_i are the executed actions according to π , which *might* be observed depending on the observation inputs. Plan recognition typically identifies π and intent recognition typically identifies G or some a_{i+j} where $j \geq 1$. A , D , and I are not shaded nodes because they are assumed, rather than observed.

ceives $o_{i+j+1}, \dots, o_{i+j+k} \in \mathcal{O}$. This continues until R_{ed} has satisfied G .

Given a subsequence of \mathcal{O} over the duration of a single action, activity recognition with simple-action hypotheses $\mathcal{H} = A$ can be performed in a similar fashion to early works in speech recognition using HMMs (Gales and Young 2007). Specifically, \mathcal{O} and \mathcal{KB} provide enough information to infer the sequence of states that generated the observations. This sequence of states provides information about the transition matrix that defines the Markov chain, and the transition matrix is more likely associated with certain actions in A . Likewise, given a sequence of these performed actions either directly observed or obtained via lower-level activity recognition, we have a new observation sequence where each $o \in \mathcal{O}'$ also satisfies $o \in A$. Then higher-level activity recognition with complex-action or macroaction hypotheses can be performed based on the transition function between inferred states in which the observed actions are applicable.

When there are no further abstractions over the actions, or if there is a task hierarchy in \mathcal{KB} , the action sequence serves as observations \mathcal{O}'' for plan recognition. This is because \mathcal{O}'' is effectively a subsequence of π , which R_{ed} is using to solve G . A is assumed to be included in \mathcal{KB} , and R_{ing} had some degree of knowledge about the world state when it began observing so that $I \in \mathcal{KB}$ as well. Then we only need to define \mathcal{H} based on the environment's context and possible plans.

If we instead define \mathcal{H} to be sets of possible goal conditions (ideally including G), then we are performing intent

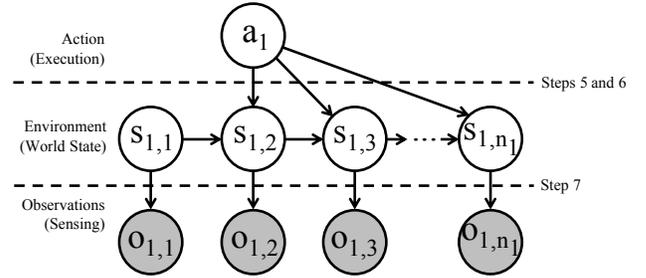


Figure 2: A focus on a single action execution HMM from Figure 1, broken down into its steps. Activity recognition typically identifies a_1 given $o_{1,1}, \dots, o_{1,n_1} \in \mathcal{O}$. These observations can range from raw signal data to discrete actions.

recognition to identify the motives behind why R_{ed} computed π in the first place. If we do not have enough observations for such an all-encompassing recognition problem, then we can also try to predict upcoming actions in π that come after the last action in \mathcal{O}'' using intent recognition.

From this perspective of reverse-engineering the plan generation and execution steps, all the types of recognition in PAIR rely on each other in order to completely understand the observed situation. An illustration of these interdependencies with respect to the HMM, higher-order latent variables, and step-by-step procedure are shown in Figures 1 and 2.

4 Using PAIR Algorithms in New Ways

The connections between each type of recognition in PAIR present new possibilities for using PAIR algorithms in the existing literature. For example, the unifying perspective above introduces the potential to daisy-chain multiple PAIR algorithms into a pipeline. If the observations are raw sensor data \mathcal{O}_{raw} , but we are interested in plan recognition using an observed sequence of actions, then we can use a lower-level activity recognition algorithm to solve each $\mathcal{R}_{activity}^i = (\mathcal{KB}_{sensor}, \mathcal{O}_{raw, start(i)...end(i)}, \mathcal{H}_{action})$ and create a new problem $\mathcal{R}_{plan} = (\mathcal{KB}_{domain}, \mathcal{O}_{action}, \mathcal{H}_{plan})$ where $o_i \in \mathcal{O}_{action}$ is derived from the solution to $\mathcal{R}_{activity}^i$. Amado et al. (2018) have a similar composition of approaches in their goal recognition algorithm for still image inputs based on latent space planning’s (Asai and Fukunaga 2018) recognition of images as grounded symbols and recognition as planning approaches. Likewise, pipelining multiple levels of activity recognition algorithms resembles the hierarchical variations of HMM used for higher-level activity recognition (Bui, Phung, and Venkatesh 2004).

We can also reconsider Ramírez and Geffner’s (2010) original probabilistic recognition as planning method to recognize plans rather than just goal conditions. Under the context that plan and intent recognition can both use the same knowledge base \mathcal{KB} and observation sequence \mathcal{O} , we only need to revise how \mathcal{KB} evaluates \mathcal{O} for the different hypotheses over plans \mathcal{H}_{plan} rather than goal conditions \mathcal{H}_{goal} . These two sets of hypotheses are also related to each other because the goal constrains which plans will be successful and vice-versa. This is emphasized by their equation for estimating the probability of the observations given a goal:

$$P(\mathcal{O} | G) = \sum_{\pi \in \Pi_G} P(\mathcal{O} | \pi, G) \cdot P(\pi | G)$$

where Π_G is the set of plans that solve G . To avoid an infinite sum, they assume that the most optimal plan is a sufficient replacement for Π_G , and we can make a closer approximation (Freedman et al. 2018) using the slightly larger set $\mathcal{H}_{plan} \cap \Pi_G$.

The probability is derived as the likelihood from Bayes’s Rule on $P(G | \mathcal{O})$, where $G \in \mathcal{H}_{goal}$. Thus we now compute $P(\pi | \mathcal{O})$; Bayes’s Rule derives the likelihood:

$$P(\mathcal{O} | \pi) = \sum_{G \in \mathcal{H}_{goal}} P(\mathcal{O} | \pi, G) \cdot P(G | \pi).$$

The shared term $P(\mathcal{O} | \pi, G)$ is not too useful because it is binary (1 if \mathcal{O} is a subsequence of π and 0 otherwise) by Ramírez and Geffner’s definition, but $P(G | \pi)$ can be found using Bayes’s Rule and the assigned probability

$$P(\pi | G) = e^{-\beta c(\pi)}$$

for some constant β . Specifically,

$$P(G | \pi) = Z^{-1} P(\pi | G) \cdot P(G)$$

where Z is a normalizing constant summing over \mathcal{H}_{goal} . As a prior over \mathcal{H}_{goal} is contained in \mathcal{KB} for the original algorithm, this alternative probability can be computed.

Geib et al. (2016) reveal another example that integrates plan and intent recognition without explicitly mentioning it. As the observations are done in real time to later negotiate how R_{ing} may assist, the derived explanations from the ELEXIR algorithm include predictions for upcoming actions. This predictive case of intent recognition is used to both ask R_{ed} which of the predicted tasks R_{ing} can do and to avoid directly getting in R_{ed} ’s way.

5 Discussion

Research within the PAIR community began with the singular notion of understanding agents based on the decisions they made, and has since divided the task into three primary types of recognition. The community still collaborates and acknowledges that they share common research questions; however, the literature does not often consider how research in one type of recognition can apply to others. To facilitate and encourage finding these relations, we introduce a formal definition for recognition as a whole that abstracts problems to a knowledge base, observation sequence, and set of hypotheses. Though the knowledge base is the most distinguishing feature between problems even within the same type of recognition, we believe that it may lead to identifying characteristics of various classes of recognition problems.

In addition to a formal definition of recognition, we propose a perspective on the overall recognition process based on reverse-engineering how an agent makes decisions and consequently acts during observation. This context unifies how each type of recognition plays a role in the fundamental goal of PAIR research, and we provide some examples based on existing recognition techniques that emphasize how they work together. Although no one formalization and perspective is usually sufficient to describe an entire field of research, we hope this will serve as an initial step towards exploring what defines PAIR and how all types of recognition are related.

Acknowledgements

We thank Sarah Keren for her discussions with the authors regarding the formalization, and we additionally thank the reviewers for their useful feedback.

References

- Amado, L.; Pereira, R. F.; Aires, J. P.; Magnaguagno, M. C.; Granada, R.; and Meneguzzi, F. 2018. Goal recognition in latent space. In *Proceedings of the 2018 International Joint Conference on Neural Networks*, 1–8.
- Anjum, A., and Ilyas, M. U. 2013. Activity recognition using smartphone sensors. In *Proceedings of the IEEE Consumer Communications and Networking Conference*, 914–919.
- Asai, M., and Fukunaga, A. 2018. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Bäckström, C. 1992. Equivalence and tractability results for SAS+ planning. In *Proceedings of the Third International*

- Conference on Principles of Knowledge Representation and Reasoning*, 126–137.
- Baker, C. L.; Saxe, R.; and Tenenbaum, J. B. 2009. Action understanding as inverse planning. *Cognition* 113(3):329–349. Reinforcement learning and higher cognition.
- Bellman, R. E. 2003. *Dynamic Programming*. Dover Books on Computer Science Series. Mineola, New York, USA: Dover Publications.
- Bui, H. H.; Phung, D. Q.; and Venkatesh, S. 2004. Hierarchical hidden Markov models with general state hierarchy. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 324–329.
- Bui, H. H.; Venkatesh, S.; and West, G. 2002. Policy recognition in the abstract hidden Markov model. *Journal Artificial Intelligence Research* 17(1):451–499.
- Callaway, F.; Lieder, F.; Krueger, P. M.; and Griffiths, T. L. 2017. Mouselab-MDP: A new paradigm for tracing how people plan. In *Proceedings of the Third Multidisciplinary Conference on Reinforcement Learning and Decision Making*, 335–339.
- Charniak, E., and Goldman, R. 1991. Probabilistic abduction for plan recognition. Technical Report CS-91-12, Brown University, Providence, RI, USA.
- Chen, Y.; Diethe, T.; and Flach, P. A. 2016. AdlTM: A topic model for discovery of activities of daily living in a smart home. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 1404–1410.
- Chikhaoui, B.; Wang, S.; and Pigot, H. 2012. ADR-SPLDA: Activity discovery and recognition by combining sequential patterns and latent dirichlet allocation. *Pervasive and Mobile Computing Special Issue on Pervasive Healthcare* 8(6):845–862.
- Cimatti, A.; Pistore, M.; Roveri, M.; and Traverso, P. 2003. Weak, strong, and strong cyclic planning via symbolic model checking. *Artificial Intelligence* 147(1-2):35–84.
- E-Martín, Y., and Smith, D. E. 2017. Goal recognition with noisy observations. In *Proceedings of the AAAI Workshop on Plan, Activity, and Intent Recognition*, 804–811.
- E-Martín, Y.; R-Moreno, M. D.; and Smith, D. E. 2015. A fast goal recognition technique based on interaction estimates. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 761–768.
- Erol, K.; Hendler, J.; and Nau, D. S. 1994. HTN planning: Complexity and expressivity. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1123–1128.
- Fine, S.; Singer, Y.; and Tishby, N. 1998. The hierarchical hidden Markov model: Analysis and applications. *Machine Learning* 32(1):41–62.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable planning. In *IJCAI-17 Workshop on Explainable AI (XAI) Proceedings*, 24–30.
- Freedman, R. G., and Zilberstein, S. 2017. Integration of planning with recognition for responsive interaction using classical planners. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 4581–4588.
- Freedman, R. G.; Fung, Y. R.; Ganchin, R.; and Zilberstein, S. 2018. Towards quicker probabilistic recognition with multiple goal heuristic search. In *Proceedings of the AAAI Workshops: Plan, Activity, Intent Recognition (PAIR)*, In Press.
- Freedman, R. G.; Jung, H.-T.; and Zilberstein, S. 2014. Plan and activity recognition from a topic modeling perspective. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling*, 360–364.
- Fu, J.; Calderon Jaramillo, A.; Ng, V.; Bastani, F. B.; and Yen, I.-L. 2016. Fast strong planning for fully observable nondeterministic planning problems. *Annals of Mathematics and Artificial Intelligence* 78(2):131–155.
- Gales, M., and Young, S. 2007. The application of hidden Markov models in speech recognition. *Foundations and Trends in Signal Processing* 1(3):195–304.
- Geib, C. W., and Goldman, R. P. 2009. A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence* 173(11):1101–1132.
- Geib, C., and Steedman, M. 2007. On natural language processing and plan recognition. In *IJCAI*, 1612–1617.
- Geib, C.; Weerasinghe, J.; Matskevich, S.; Kantharaju, P.; Craensen, B.; and Petrick, R. P. A. 2016. Building helpful virtual agents using plan recognition and planning. In *Twelfth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 162–168.
- Geib, C. W. 2009. Delaying commitment in plan recognition using combinatorial categorial grammars. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI’09, 1702–1707. Pasadena, California, USA: Morgan Kaufmann Publishers Inc.
- Gerevini, A. E.; Haslum, P.; Long, D.; Saetti, A.; and Dimopoulos, Y. 2009. Deterministic planning in the fifth international planning competition: {PDDL3} and experimental evaluation of the planners. *Artificial Intelligence* 173(5–6):619–668. Advances in Automated Plan Generation.
- Goldman, R. P.; Geib, C. W.; Kautz, H.; and Asfour, T. 2011. Plan Recognition – Dagstuhl Seminar 11141. *Dagstuhl Reports* 1(4):1–22.
- Gori, I.; Sinapov, J.; Khante, P.; Stone, P.; and Aggarwal, J. 2015. Robot-centric activity recognition ‘in the wild’. In *Proceedings of the International Conference on Social Robotics (ICSR)*, 224–234.
- Harmon, L. D. 1982. Automated tactile sensing. *The International Journal of Robotics Research* 1(2):3–32.
- Hobbs, J. R.; Stickel, M. E.; Appelt, D. E.; and Martin, P. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1):69–142.
- Huỳnh, T.; Fritz, M.; and Schiele, B. 2008. Discovery of activity patterns using topic models. In *Proceedings of the Tenth International Conference on Ubiquitous Computing*, 10–19.
- Inoue, N., and Inui, K. 2011. IIP-based reasoning for weighted abduction. In *Proceedings of the AAAI Workshops: Plan, Activity, Intent Recognition (PAIR)*, 25–32.

- Jung, H.-T.; Freedman, R. G.; Foster, T.; Choe, Y.-K.; Zilberstein, S.; and Grupen, R. A. 2015. Learning therapy strategies from demonstration using latent Dirichlet allocation. In *Proceedings of the Twentieth International Conference on Intelligent User Interfaces*, 432–436.
- Kautz, H. A. 1987. *A formal theory of plan recognition*. Ph.D. Dissertation, University of Rochester.
- Kautz, H. A. 1991. A formal theory of plan recognition and its implementation. In Allen, J. F.; Kautz, H. A.; Pelavin, R. N.; and Tenenber, J. D., eds., *Reasoning About Plans*. San Francisco, CA, USA: Morgan Kaufmann. 69–125.
- Keren, S.; Gal, A.; and Karpas, E. 2014. Goal recognition design. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling*, 154–162.
- Keren, S.; Gal, A.; and Karpas, E. 2015. Goal recognition design for non-optimal agents. In *Proceedings of the Twenty-Ninth International Conference on Artificial Intelligence*, 3298–3304.
- Levine, S., and Williams, B. 2014. Concurrent plan recognition and execution for human-robot teams. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling*, 490–498.
- Madani, O.; Bui, H. H.; and Yeh, E. 2009. Efficient online learning and prediction of users’ desktop actions. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 1457–1462.
- Mainprice, J.; Hayne, R.; and Berenson, D. 2016. Goal set inverse optimal control and iterative replanning for predicting human reaching motions in shared workspaces. *IEEE Transactions on Robotics* 32(4):897–908.
- Mirsky, R., and Gal, Y. K. 2016. SLIM: Semi-lazy inference mechanism for plan recognition. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 394–400.
- Mirsky, R.; Gal, Y. K.; and Shieber, S. M. 2017. CRA-DLE: an online plan recognition algorithm for exploratory domains. *ACM Transactions on Intelligent Systems and Technology* 8(3):45:1–45:22.
- Pednault, E. P. D. 1989. ADL: Exploring the middle ground between STRIPS and the situation calculus. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, 324–332. Toronto, Canada: Morgan Kaufmann Publishers Inc.
- Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-based heuristics for goal recognition. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 3622–3628.
- Ramírez, M., and Geffner, H. 2009. Plan recognition as planning. In *Proceedings of the Twenty-first International Joint Conference on Artificial Intelligence*, 1778–1783.
- Ramírez, M., and Geffner, H. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 1121–1126.
- Ramírez, M., and Geffner, H. 2011. Goal recognition over POMDPs: Inferring the intention of a POMDP agent. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 2009–2014.
- Riedl, M. O., and Young, R. M. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39(1):217–268.
- Rieping, K.; Englebienne, G.; and Kröse, B. 2014. Behavior analysis of elderly using topic models. *Pervasive and Mobile Computing* 15(0):181–199.
- Sanner, S. 2010. Relational dynamic influence diagram language (RDDL): Language description. *Technical Report, Australian National University*.
- Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan recognition as planning revisited. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 3258–3264.
- Sukthankar, G.; Geib, C.; Bui, H. H.; Pynadath, D.; and Goldman, R. P. 2014. *Plan, Activity, and Intent Recognition: Theory and Practice*. Waltham, MA, USA: Elsevier Science.
- Thorpe, C.; Hebert, M. H.; Kanade, T.; and Shafer, S. A. 1988. Vision and navigation for the carnegie-mellon navlab. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10(3):362–373.
- Unhelkar, V. V.; Pérez-D’Arpino, C.; Stirling, L.; and Shah, J. A. 2015. Human-robot co-navigation using anticipatory indicators of human walking motion. In *IEEE International Conference on Robotics and Automation*, 6183–6190.
- Vered, M., and Kaminka, G. A. 2017. Heuristic online goal recognition in continuous domains. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 4447–4454.
- Wayllace, C.; Hou, P.; Yeoh, W.; and Son, T. C. 2016. Goal recognition design with stochastic agent action outcomes. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, 3279–3285. New York, New York, USA: AAAI Press.
- Younes, H. L. S., and Littman, M. L. 2004. PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects. Technical Report CMU-CS-04-162, Carnegie Mellon University, Pittsburgh, PA, USA.
- Zhang, H., and Parker, L. E. 2011. 4-dimensional local spatio-temporal features for human activity recognition. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2044–2049.