

Session Analysis using Plan Recognition

Reuth Mirsky and Ya'akov (Kobi) Gal

Department of Software and Information Systems Engineering
Ben-Gurion University of the Negev, Israel
{dekelr,kobig}@bgu.ac.il

David Tolpin

PayPal, Israel
dtolpin@paypal.com

Abstract

This paper presents preliminary results of our work with a major financial company, where we try to use methods of plan recognition in order to investigate the interactions of a customer with the company's online interface. In this paper, we present the first steps of integrating a plan recognition algorithm in a real-world application for detecting and analyzing the interactions of a customer. It uses a novel approach for plan recognition from bare-bone UI data, which reasons about the plan library at the lowest recognition level in order to define the relevancy of actions in our domain, and then uses it to perform plan recognition.

We present preliminary results of inference on three different use-cases modeled by domain experts from the company, and show that this approach manages to decrease the overload of information required from an analyst to evaluate a customer's session — whether this is a malicious or benign session, whether the intended tasks were completed, and if not — what actions are expected next.

1 Introduction

Online websites are often open-ended and flexible interfaces, supporting interaction styles by users that include exogenous actions and high noise rate. Such interfaces provide a rich framework for users and are becoming increasingly prevalent in many web and mobile applications, but challenge conventional plan recognition algorithms for inferring users' activities with the software.

The open-ended nature of these settings afford a rich spectrum of interaction for clients: they can perform the same task in many different ways (such as logging in through a website or by a mobile application), engage in exploratory activities involving trial-and-error, they can repeat activities indefinitely (browsing), and they can interleave between activities.

This paper focuses on inferring customers' activities for an Internet financial company, in which customers widely engage in exploratory behavior, and present initial results for plan recognition in such settings that can be the basis of future work in this direction.

Our data set consists of customer sessions of a large

financial company¹. The customer sessions are recorded as low-level activities of 'click stream' — web page visits and user interface actions. This data is a rich source of information about customer's behavior, but cannot be used as-is to describe the intentions or plans of the customer: buy product, sell product, manage account, etc. The underlying model used for the recognition process is a set of hierarchical task networks (HTNs) called a "plan library", that allows hierarchical decomposition of a task, which can be used both as an explanatory accessory to an overseer and for prediction of future actions. Our main goal is to detect frauds, validate transactions and predict future actions — both in terms of the user's intents and in terms of the execution of these intents.

The basic stream of information can be thought of as a recording of customer's *conversation* with the system. Understanding the language of conversation would bring benefits in various areas. For example:

- Predictive verification — by following customer's actions in real time, we can have our systems guess next actions and alert when an agent biases from their predictive plan.
- Proactive validation — if we can understand the purpose of customer's visit based on the beginning of their activity, we can let our systems start validation of the session even before it takes place, facilitating faster response and better utilization of resources.
- Computer security — certain types of fraud are performed by malicious software controlling customer computers. We may be able to distinguish between malicious and benign actors on the other side of the wire.

However, as a flexible interface, the system allows the user to engage in more than one task at a time (making multiple purchases), make mistakes (pressing a button twice instead of once), repeat actions (browse between different parts of the interface) and more. All of these are features that define an exploratory environment (GRS+12).

¹All interactions and examples were altered to preserve privacy

While these traits make the interface useful and convenient to the user, it challenges conventional plan recognition algorithms for inferring users’ activities with the software, as they need to reason about all these factors (AG13; UDSG15).

2 Related Work

The tasks described in the introduction are traditionally addressed by activity recognition algorithms, which are used to tokenize the stream of data into detectable actions (BI04; LFK06; FZ17). However, these algorithms usually cannot perform higher levels of inference to describe the agent’s behavior or predict the future actions of the agent (HHP16; Yor17). Different approaches also used an HMM representation or data-driven learning to elicit tasks from low-level activities (NBT⁺08; MBY09).

Some notable exceptions are works by (QL04; DBPV05; TBC⁺14) which try to combine low level activities and higher level domain knowledge, or (TBC⁺14) which propose a multi-agent model for robot collaboration based on plan recognition. However, the works in these lines of research tend to use a domain-theory based recognition or other probabilistic representations, which do not capture the hierarchical nature of task decomposition like Hierarchical task networks (HTNs) and plan libraries.

In order to address the above tasks, we need to use hierarchical plan recognition, a field of research exploring algorithms that recognize the plans of the agent based on a partial sequence of actions, and predict future actions (Kau87). Few works (GPM⁺17; GG01) did use a plan library as the underlying domain knowledge for the task, but they did not provide predictions or used this information to formalize the output. Other works (JLM05) do output a complete hierarchy, but the plan recognition algorithm used does not work well in exploratory environments.

In this work, we present a criteria for constructing the plan library that will then be used by the agent to perform the plan recognition task. Thus, we do not need to use an activity recognizer. We integrate instead the low-level activities directly into the plan recognition task.

3 Definitions

Plan recognition takes as input a plan library and a partial sequence of observations, and outputs a set of possible explanations.

3.1 Plan Library

The plan library can be thought of as the language of actions, as well as goals we want to be able to recognize. The library must specify *terminals*, *non-terminals*, *goals*, and *derivation rules*.

Definition 1 (*Explanation*) *A plan library is a tuple $PL = \langle T, NT, G, R \rangle$, where T is a set of are low-level, observable action, NT is a set of complex level actions,*

composed from either T s or other NT s, G is a subset of non-terminals corresponding to the highest level of actions, representing the goals the agent can achieve and R is a set of derivation rules which specify how each complex action can be decomposed into other actions of the form $\langle nt \rangle \rightarrow \langle sequence \rangle \mid \langle order \rangle$, where:

- $\langle nt \rangle$ is a non-terminal, a complex action.
- $\langle sequence \rangle$ is a sequence of actions which compose the complex action.
- $\langle order \rangle$ is a list of tuples defining the the ordering of actions — actions must appear in the same order as in the tuple, but ordering between different tuples is unspecified. For example, ‘(login, addName)’ states that the user must log-in before adding a card.

Consider the following very simple plan library for our domain.

- $T = \{login, addName, addCredit, signup, submit, home, payment, success, transfer, confirm\}$
- $NT = \{AddAccount, Buy\}$
- $G = \{AddAccount, Buy\}$
- $R =$

$AddAccount$	\rightarrow	$login, addName, addCredit$	
		$[(login, addName)(addName, addCredit)]$	
$AddAccount$	\rightarrow	$signup, addName, submit$	
		$[(signup, addName)(addName, submit)]$	
Buy	\rightarrow	$home, payment, success$	
		$[(home, payment)(payment, success)]$	
Buy	\rightarrow	$home, transfer, confirm$	
		$[(home, transfer)(transfer, confirm)]$	

3.2 Plans and Explanations

Based on the definitions above, we can apply an algorithm to a sequence of actions to obtain explanations about ongoing behaviors and anticipated future actions.

A plan is a labeled tree $p = (V, E, \mathcal{L})$, where V and E are the nodes and edges of the tree, respectively, and \mathcal{L} is a labeling function $\mathcal{L} : V \rightarrow NT \cup T$ mapping every node in the tree to either a basic or a complex action in the plan library. Each inner node is labeled with a complex action such that its children nodes are a decomposition of its complex action into constituent actions according to one of the refinement methods. The set of all leaves of a plan p is denoted by $leaves(p)$, and a plan is said to be *complete* iff all its leaf nodes are labeled basic actions, i.e., $\forall v \in leaves(p), \mathcal{L}(v) \in T$.

An *observation sequence* is an ordered set of basic actions that represents actions carried out by the observed agent. A plan p *describes* an observation sequence O iff every observation is mapped to a leaf in the tree. Formally, there exists an injective function $f : O \rightarrow leaves(p) \cap T$ such that $f(o) = v$. The observed agent is assumed to plan by choosing a subset of complex actions as intended goals and then carrying out a separate plan for completing each of these goals.

An agent may pursue several goals at the same time. Therefore, an explanation can include a set of plans, as described in the following definition:




Figure 1: An Explanation with Three Plans.

Definition 2 (*Explanation*) An explanation for an observation sequence is a set of plans such that each plan describes a mutually exclusive subset of the observation sequence and taken together the plans describe all of the observations. We then say that the explanation describes the observation sequence.

For the rest of this paper we will use a running example based on the following sequence of processed observations: $[home, login, addName, login, addCredit]$ When we submitted the sequence and the plan library to a plan recognition algorithm, we obtained two explanations. The first one is presented in Figure 1.

This explanation is a set of trees. In each tree the leaves are either the actions we observed, or actions predicted to perform in the future. A special tag *frontier* specifies the location in the tree from where the expansion may continue, based on future actions. Leaves with this tags are colored in green in the figure. Future actions that still has preceding actions that have not been executed are colored with blue. If a tree has no ‘frontier’ leaf (top tree in our example), then the goal is fully described by actions.

The output above represents the following explanation of the input sequence: (1) The agent performed one buy transaction and two account additions; (2) The agent performed the first action in the buy task, but then stopped – we expect that the next action to complete this task is *transfer*; (3) Regarding the account additions, one was completed and the other was not.

For our very basic example, one can think about a brute-force algorithm for explaining the activity. However in general, a sophisticated algorithm is required to provide short and meaningful explanations quickly, in a manner suitable for online processing. The remaining of this paper talks about the very beginning of work in collaboration with analysts and developers from the financial company.

4 Recognition Components

We now detail the components we use to perform the plan recognition task from the raw click stream of customer sessions.

4.1 Preprocessing

A session is the basic unit of interaction between a customer and the system. Efficient analysis and validation of sessions in real time is crucial for technical realization of the company’s business. The raw input in our domain is a stream of customer sessions. Each session is constructed from a list of entries, where each entry has a specific timestamp, user information and a special label describing the page on the company’s site the user is visiting. The average length of a session was 80.49 entries.

As a first step, we needed to reduce this number as many entries are irrelevant to the customer’s main tasks in the system and integrating them into terminals in the plan library would be both inefficient and less informative. We defined the following criteria for a session entry to be considered relevant to us:

Definition 3 (*Landmark in PL*) A raw entry e in a complete session S is considered a landmark in relation to a plan library PL if $S \setminus \{e\}$ cannot describe a plan from PL .

In collaboration with our partners inside the company, we elicited about 20 types of entries, defined by their page labels, which are landmarks for executing the relevant tasks. We classified each of the pages into a suitable basic action as they appear in our plan library. Counting only these pages as relevant actions we wish to observe, we received an average of 9.68 observations per session (with stdev of 7.26).

However, as an exploratory environment, many entries in a session are redundant even if they can be considered landmarks, as the user can make mistakes or perform a landmark action more times than necessary. To perform an elimination of these entries and to remain only with the most relevant observations, we performed another sifting. We discarded all observations that were a repetition of the earlier observation. After this elimination, we reduced the average number of relevant observations per session to 3.68 (stdev=2.62).

4.2 Recognition

There are many possible strategies that a customer can use to perform interactions, and variations within each due to exploratory activities and mistakes carried out by the customer. Even after the preprocessing we enforced on the set of actions performed by the customer, there are still two types of exploratory behaviors which can hinder the ability of the recognizer to infer the customer’s actions correctly: (1) Exogenous actions: even after our preprocessing effort, many actions cannot be combined together when looking at the complete sequence, thus the plan recognizer cannot output any explanation that described all of the actions in the session. (2) One explanation may relate a given action to a relevant task, while another may relate this action to a failed attempt or a mistake. The space of possible explanations can become very large, even for a small number of observations.

The CRADLE algorithm (MGS17) proposes solutions for both of these problems using three components:

Inference it receives as input a plan library, an observation sequence, and outputs a set of explanations, each of which is an explanation of the observation sequence in the sense of Definition 2.

Filters it filters redundant explanations according to a set of domain-independent conditions. A filter is a function taking a candidate explanation e , and returning *true* or *false* depending on whether the candidate explanation does or does not pass a certain condition. For our presented domain, we tried several values and filter types and finally set 3 filters: (1) The number of plans in the explanation is less than or equal to the average; (2) The number of frontier nodes in the explanation is less than or equal to the average; (3) The number of different plans in the explanation is less than 4. We discard all explanations which do not pass these thresholds

Exogenous Actions CRADLE can handle exogenous actions and mistakes and can omit them from the set of explanations as needed.

5 Empirical

We used a click stream of selected sessions, labeled by the company’s analysts as sessions containing the relevant tasks. In total, we tested 3 types of sessions, with 50 instances of each session type: Buy, Add account for existing user (AAExist) and Add account for new user (AANew). For each of these sessions, we performed the process described in Section 4.

Figure 2 presents the average number of different types of explanations outputted by CRADLE per task: Total is the total number of explanations; Full Plan is the number of explanations in which at least on plan was completed; No Open means the number of explanations in which all plans were completed and had no open frontier. Notice that even in sessions that were labeled by analysts as relevant to some task, the plan recognition process shows that only a small portion of the sessions contain a completed task. These results are similar in the PHATT runs, even though it does not discard explanations, thus might output more explanations with a completed task. We intend to validate these results with other data to evaluate if these sessions were indeed sessions of incomplete tasks.

In order to evaluate the performance of CRADLE, we compared the runtime and number of outputted explanation in comparison to the PHATT algorithm, augmented with the exogenous actions handling of CRADLE (without this augmentation, PHATT would not be able to output any explanation at all that describes the complete sequence of observations). Table 5 summarizes these runs.

The first and most important thing to notice in the table, is the difference between the original length of a session entry, and the outputted set of explanations




Figure 2: Average Number of Explanations by Type.

by the plan recognizers. An average decrease of 83% in the number of entities representing the session. Such a decrease allows a faster analysis of the session, since the number of possible explanations is exponential with the number of observations (MGS17). For example, even with the pruning of CRADLE, the average number of explanations it outputted for the AANew sessions is 146.06 (compared to 1.33 with the preprocessing). Moreover, the outputted explanations have a structure imposed by the plan library that the original stream lacked.

The second point to notice it that the values of the AAExist case are similar in both algorithms. We attribute this to the fact that performing this task is always performed in the same fashion, without any options to bias in the interface of the company’s website.

6 Future Work

This a preliminary work for using plan recognition for exploratory environments on real-world click streams. It uses a novel approach for plan recognition from bare-bone UI data, which reasons about the plan library in the lowest recognition level in order to define the relevancy of actions in our domain, and then uses it to perform plan recognition.

While we manage to process low level data using mostly tools intended for higher levels of inference, there is still much to be done from both ends of the process: First, we expect that we can use most intelligent tools for the preprocessing stage, either from the world of activity recognition or natural language pro-

	Buy	AAExist	AANew
Session Entries	147.02	14.13	80.33
Observations	5.64	2.06	3.33
CRADLE Explanations	8.74	1.00	1.33
PHATT Explanations	18.32	1.00	2.00
CRADLE Time (seconds)	0.06	0.01	0.01
PHATT Time (seconds)	0.07	0.01	0.03

Table 1: Runtime and Explanation Set Size for CRADLE and PHATT

cessing. As we started to show here, we wish to use the domain knowledge (e.g. the plan library) in our low-level action detection. Second, we wish to visualize this information in a coherent manner that will allow analysts evaluate sessions in real time for fast verification and validation, or to be able to counter adversarial behavior in time.

References

- [AG13] O. Amir and Y. Gal. Plan recognition and visualization in exploratory learning environments. *ACM Transactions on Interactive Intelligent Systems*, 3(3):16:1–23, 2013.
- [BI04] Ling Bao and Stephen S Intille. Activity recognition from user-annotated acceleration data. In *International Conference on Pervasive Computing*, pages 1–17. Springer, 2004.
- [DBPV05] T. V. Duong, H. H. Bui, D. Q. Phung, and S. Venkatesh. Activity recognition and abnormality detection with the switching hidden semi-markov model. In *Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 838–845. IEEE, 2005.
- [FZ17] Richard G. Freedman and Shlomo Zilberstein. Integration of planning with recognition for responsive interaction using classical planners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017.
- [GG01] Christopher W Geib and Robert P Goldman. Plan recognition in intrusion detection systems. In *DARPA Information Survivability Conference & Exposition II, 2001. DISCEX'01. Proceedings*, volume 1, pages 46–55. IEEE, 2001.
- [GPM⁺17] Roger Granada, Ramon Fraga Pereira, Juarez Monteiro, Duncan D Ruiz, Rodrigo Coelho Barros, and Felipe Meneguzzi. Hybrid activity and plan recognition for video streams. In *Proceedings of the AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*, 2017.
- [GRS⁺12] Y. Gal, S. Reddy, S. Shieber, A. Rubin, and B. Grosz. Plan recognition in exploratory domains. *Artificial Intelligence*, 176(1):2270 – 2290, 2012.
- [HHP16] Nils Y Hammerla, Shane Halloran, and Thomas Ploetz. Deep, convolutional, and recurrent models for human activity recognition using wearables. *arXiv preprint arXiv:1604.08880*, 2016.
- [JLM05] Peter A Jarvis, Teresa F Lunt, and Karen L Myers. Identifying terrorist activity with ai plan recognition technology. *AI Magazine*, 26(3):73, 2005.
- [Kau87] H. A. Kautz. *A formal theory of plan recognition*. PhD thesis, University of Rochester, 1987.
- [LFK06] Lin Liao, Dieter Fox, and Henry Kautz. Location-based activity recognition. *Advances in Neural Information Processing Systems*, 18:787, 2006.
- [MBY09] Omid Madani, Hung Hai Bui, and Eric Yeh. Prediction and discovery of users’ desktop behavior. In *AAAI Spring Symposium: Human Behavior Modeling*, pages 49–55, 2009.
- [MGS17] Reuth Mirsky, Ya’akov (Kobi) Gal, and Stuart Shieber. Cradle: An online plan recognition algorithm for exploratory domains. In *ACM Transactions on Intelligent Systems and Technology*, 2017.
- [NBT⁺08] Sriraam Natarajan, Hung H Bui, Prasad Tadepalli, Kristian Kersting, and Weng-Keen Wong. Logical hierarchical hidden markov models for modeling user activities. In *International Conference on Inductive Logic Programming*, pages 192–209. Springer, 2008.
- [QL04] Xinzhou Qin and Wenke Lee. Attack plan recognition and prediction using causal networks. In *Computer Security Applications Conference, 2004. 20th Annual*, pages 370–379. IEEE, 2004.
- [TBC⁺14] Kartik Talamadupula, Gordon Briggs, Tathagata Chakraborti, Matthias Scheutz, and Subbarao Kambhampati. Coordination in human-robot teams using mental modeling and plan recognition. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 2957–2962. IEEE, 2014.
- [UDSG15] Oriel Uzan, Reuth Dekel, Or Seri, and Ya’akov Gal. Plan recognition for exploratory learning environments using interleaved temporal search. *AI Magazine*, 36(2):10–21, 2015.
- [Yor17] Kristina Yordanova. Texttohbm: A generalised approach to learning models of human behaviour for activity recognition from textual instructions for activity recognition from textual instructions. In *Proceedings of the AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*, 2017.