Plan Recognition in Exploratory Domains (Doctoral Consortium)

Reuth Mirsky
Supervisor: Ya'akov (Kobi) Gal
Dept. of Information Systems Engineering
Ben-Gurion University
Beer-Sheva 84105, Israel
dekelr@post.bgu.ac.il

ABSTRACT

My thesis focuses on recognition and interventions of users' activities using open-ended and flexible software. In such *exploratory settings*, users' behavior is characterized by exploration, mistakes and trial-and-error. Exploratory domains provide a flexible and rich interaction environment for their users, but induce challenges for automatic recognition and support of their activities. My thesis focuses on the following three challenges which are central to understanding users' interactions in exploratory settings and to use this understanding in order to provide them with effective support and guidance: (1) Representing and inferring users' interactions in exploratory domains. (2) Disambiguating between possible explanations in order to improve understanding of users' behavior. (3) Producing machine-generated support that adapts to the needs of the users. My research activities combines computational models, algorithms and empirical methodologies to meet the challenges above. They are conducted in the context of various types of exploratory settings. Specifically, I am developing novel plan recognition algorithms for inferring users' interactions in exploratory settings and intervention mechanisms for these environments. I am evaluating my approach in the real world using educational software, medical records and cyber security domains. My results so far include (1) design of a new model for plan recognition; (2) an online plan recognition algorithm that is empirically shown to outperform the state-of-the-art methods in the real world; (3) A sequential process that allows informed disambiguation of possible hypotheses describing an agent's plans. The long term impact of my contribution to computer science will be demonstrated by (1) developing new algorithms for plan recognition, intervention design and adaptation for exploratory settings; (2) showing that these methods generalize to different types of settings that differ widely in they type of interaction that is provided by the users.

Keywords

Human-Aware AI; Activity and Plan Recognition; Model-Based Reasoning; Exploratory Settings

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

1. INTRODUCTION

The focus of my thesis is on algorithms and techniques for inferring and guiding agents' behavior in domains in which users engage in exploratory and error-prone behaviors. Developing technology is changing rote and monolithic interaction styles between computers and their users are broken to more flexible types of interactions that allow users to explore and interleave between different activities. Examples include Integrated Development Environments (IDEs), security monitoring software [1], and educational software [9].

One problem I addressed in my thesis online plan recognition in real world exploratory domains [8]. I defined a model for agent's behavior in exploratory domains that is based on a context sensitive grammar with parameterization and free ordering. I augmented an existing state of the art online plan recognition algorithm to be able to consider such domains. I devised several heuristics to make the algorithm run on real-world domains feasible. The heuristics I used are domain dependant, such as the number of future actions expected, and the coherency of the hypotheses. I tried these heuristics on several domains. First, I used the simulated domain generated for a different plan recognition algorithm and compared between them [6]. Second, I used real world logs of students' behavior using educational software [9]. I managed to show that both in the simulated settings and in the real-world domains, the augmented plan recognition algorithm can handle large amount of input and can also output precise explanations of the agent's actions, either simulated or students [3, 4].

Another problem I am tackling is the disambiguation problem in plan recognition. Due to imperfect knowledge about the agent's behavior and the environment, it is often the case that there are multiple hypotheses about an agent's plans that are consistent with the observations, though only one of these hypotheses is correct. My approach assumes the acting agent can be queried about whether a candidate plan in a hypotheses is correct. Only few works exist on interacting with the observed agent as means to disambiguate the hypothesis space during plan recognition [2, 5], but none explicitly reason about the best queries that will maximize the information gain. I defined the sequential plan recognition problem, which seeks to minimize the number of queries made that are required to find the correct hypothesis. In order to address this problem, I propose a process that is performed sequentially and used to update the set of possible hypotheses during the recognition process. Initial results on the same domains as the previous project show that the use of reasoning can decrease significantly the number of queries required to disambiguate between hypotheses.

This thesis will develop more representations and algorithms for building collaborative tools for supporting the participants of exploratory settings that will entail the following capabilities: (1) Automatically recognizing users' activities in real-time during their interaction with the software; (2) Reasoning about the recognized hypotheses and reducing ambiguity of algorithm's output; (3) Producing machine-generated support that adapts to the needs of all participants, including system, users and overseers.

My thesis will combine computational models, algorithms and empirical methodologies that augment existing software to provide capabilities for meeting the desiderata above. My empirical methodology will be conducted on data that is obtained from different types of settings. First, by using data that is obtained from real world exploratory domains, consisting of students' interactions with flexible educational software, traces of users' behavior with operating systems, and medical records. Second, by generating simulated data that allows to control the extent of exploratory activities and provide a benchmark of comparison for my algorithms.

The efficacy and robustness of my approach will be evaluated in several ways: First, by comparing the performance of my recognition and intervention algorithms to a gold standard. Second, by showing the generalizability of my techniques to the different domains. Third, by conducting user studies to compare different visualizations of users' work to overseers (e.g., teachers, researchers); Lastly, by assessing students' performance using the tools I develop.

2. EXPLORATORY GRAMMAR AND CRADLE

In exploratory domains, agents' activities map onto logs of behavior that include switching between activities, extraneous actions, and mistakes. Such settings are prevalent in real world applications such as interaction with open-ended software, collaborative office assistants and integrated development environments. Despite the prevalence of such settings in the real world (e.g., interaction with open-ended software, collaborative office assistants, integrated development environments), there is scarce work in formalizing and inferring agents' behavior in these settings. In my thesis, I developed a formal grammar for describing users' activities in such domains. I also devised a new top-down plan recognition algorithm called CRADLE that uses this grammar to recognize agents' interactions in exploratory domains. I compared the performance of CRADLE with state-of-the-art plan recognition algorithms in experimental settings that use real logs taken from students' interactions with a flexible, open-ended software for chemistry education. My empirical results show that CRADLE was able to output plans in an order of magnitude more quickly than the state-of-the-art without compromising correctness, as determined by domain experts. This approach can form the basis of future systems for providing real-time support to users in a growing class of interesting and challenging domains. My preliminary results of CRADLE were presented in EDM and AAMAS [3, 4].

3. SEQUENTIAL PLAN RECOGNITION

Plan recognition algorithms infer agents' plans from their observed actions. Due to imperfect knowledge about the agent's behavior and the environment, it is often the case that there are multiple hypotheses about an agent's plans that are consistent with the observations, though only one of these hypotheses is correct. I proposed a process to address the problem of how to disambiguate between hypotheses, by querying the acting agent about whether a candidate plan in a hypotheses matches its intentions. This process is performed sequentially and used to update the set of possible hypotheses during the recognition process. I defined the sequential plan recognition problem, which seeks to minimize the number of queries made that are required to find the correct hypothesis. It adapts a number of probing techniques for choosing which plan to query, such as maximal information gain and maximum likelihood. These approaches were evaluated on two domains from the literature using a well known plan recognition algorithm. The results showed that the information gain approach was able to find the correct plan using significantly fewer queries than the maximum likelihood approach as well as a baseline approach choosing random plans. My technique can inform the design of future plan recognition systems that interleave the recognition process with intelligent interventions of their users. Preliminary results are shown in AAMAS16 [7]

REFERENCES

[1] B. Auslander, K. M. Gupta, and D. W. Aha. Maritime threat detection using plan recognition. In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*, pages 249–254. IEEE, 2012.

[2] F. Bisson, F. Kabanza, A. R. Benaskeur, and H. Irandoust. Provoking opponents to facilitate the recognition of their intentions. In *AAAI*, 2011.

[3] R. Dekel and K. Gal. On-line plan recognition in exploratory learning environments. In *EDM (Workshops)*, 2014.

[4] R. Dekel and Y. Gal. Online plan recognition for exploratory domains. In *Workshop on Human-Agent Interaction Design and Models at AAMAS*, 2015.

[5] M. S. Fagundes, F. Meneguzzi, R. H. Bordini, and R. Vieira. Dealing with ambiguity in plan recognition under time constraints. In *Proceedings of AAMAS*, pages 389–396. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[6] F. Kabanza, J. Filion, A. R. Benaskeur, and H. Irandoust. Controlling the hypothesis space in probabilistic plan recognition. In *Proceedings of IJCAI*, pages 2306–2312. AAAI Press, 2013.

[7] R. Mirsky, Y. Gal, R. Stern, and M. Kalech. Sequential plan recognition. In *Proceedings of AAMAS*, 2016.

[8] O. Uzan, R. Dekel, and Y. Gal. Plan recognition for exploratory domains using interleaved temporal search. In *Proceedings of AIED*, 2013.

[9] D. Yaron, M. Karabinos, D. Lange, J. Greeno, and G. Leinhardt. The ChemCollective–Virtual Labs for Introductory Chemistry Courses. *Science*, 328(5978):584, 2010.