

AI Planning for Enterprise: Putting Theory Into Practice

Shirin Sohrabi
IBM Research
ssohrab@us.ibm.com

Abstract

In this paper, I overview a number of AI Planning applications for Enterprise and discuss a number of challenges in applying AI Planning in that setting. I will also summarize the progress made to date in addressing these challenges.

1 Introduction

Planning is a long-standing area of research within AI. Planning is the task of finding a procedural course of action for a declaratively described system to reach its goals while optimizing overall performance measures. AI Planning can help when (1) your problem can be described in a declarative way; (2) you have domain knowledge that should not be ignored; (3) pure learning techniques are difficult to use either because there is a structure of the problem that cannot be learned by training or that there is little to no training data available; (4) you want to be able to explain a particular course of action the system took; or (5) you can leverage the existing relationships between a problem similar to yours to AI Planning.

AI Planning has been applied to many applications including but not limited to robotics, manufacturing, logistics, transportation, and space¹. I have in particular investigated the relationship between AI Planning and several applications, which I highlight below:

- requirement engineering: helping analysts acquire a better understanding of the impact of high-level stakeholder preferences on low-level design decisions by modeling and reasoning with prioritization of stakeholder's goal [Liaskos *et al.*, 2010; Liaskos *et al.*, 2011];
- web service composition: helping users customize the compositions of services with respect to their preferences [Sohrabi and McIlraith, 2009; Sohrabi and McIlraith, 2010; Sohrabi, 2012];
- stream processing; helping analyst with the task of automating the composition of stream processing applications [Sohrabi *et al.*, 2013a];
- large-scale data analysis: helping users (network administrators, nurses, physicians) by orchestrating the data

analysis process automatically with a focus on two applications: early detection of health complications in critical care, and detection of anomalous behaviors of network hosts in enterprise networks. In both cases, we generate multiple hypotheses in order to reason about possibly incomplete, noisy, or inconsistent sequences of observations received from external sources [Sohrabi *et al.*, 2013b; Riabov *et al.*, 2015];

- future state projection: providing analysts with the ability to explain the given observations and based on that project multiple possible futures. Our focus was on energy domain application where the objective is to project the price of oil and volume of oil produced 15 years into the future [Sohrabi *et al.*, 2017]; and
- enterprise risk management: assisting financial organizations in identifying and managing emerging risks given observations derived from relevant news and social media [Sohrabi *et al.*, 2018b; Sohrabi *et al.*, 2019].

A common theme among the applications that I was involved with is the connection to an established relationship to AI Planning. In particular, in my work I leverage the existing relationship between diagnosis and planning [Sohrabi *et al.*, 2010], explanation generation as planning [Sohrabi *et al.*, 2011], and plan recognition as planning [Ramírez and Geffner, 2009; Ramírez and Geffner, 2010; Sohrabi *et al.*, 2016a]. Furthermore, in these applications, the domain knowledge exists but not necessary in a form that is accessible by an AI planner; hence a major effort is required to address the knowledge engineering and extraction problem. In addition, in applying AI Planning technologies we often encounter novel research problems. While, this novel problem is motivated by the application at hand, it has applicability and interest to the wider research community. In the next sections, I highlight these key research challenges and summarize the progress made to date in addressing them.

2 Relationship to Planning

The first and fundamental challenge is defining the application problem and making the correspondence to AI Planning. That is to first define the computational problem, and then reduce it to a planning problem or to a problem in which an established relationship to AI Planning exists (e.g., plan

¹<http://users.cecs.anu.edu.au/~patrik/sigaps/>

recognition). This allows us to use AI Planning tools to compute a solution to the original defined computational problem.

For example, in the case of large-scale data analysis application, we first define what we call a hypothesis generation problem, and then leverage the existing relationship between generating explanations [Sohrabi *et al.*, 2011], a more general form of diagnosis, and a planning problem with temporally extended goals. In the enterprise risk management application, we first formally define the scenario planning problem, and then make a correspondence between the scenario planning problem and the plan recognition problem. We then leverage the existing relationship between plan recognition and AI Planning [Ramírez and Geffner, 2010; Sohrabi *et al.*, 2016a] to compute the solution to the original scenario planning problem.

Note that we had to extend the established relationships in a number of directions in order to address the challenges that present themselves in such settings. For example, for plan recognition as planning, we consider observations that are over fluents (i.e., properties of states), and that they can be unreliable (i.e., not all of them are explainable).

3 Knowledge Engineering

Knowledge engineering is an important challenge when it comes to designing any system, where domain knowledge is encoded and maintained directly by the domain experts, such as clinicians, network security engineers, or analysts. To address this challenge, I worked on two separate solutions. The first one was to develop a language called LTS++ [Sohrabi *et al.*, 2016c] that allows the domain experts to easily describe the state transition models and observations specific to their domain, without requiring the experts to learn about the underlying planning technologies or Planning Domain Definition Language (PDDL) [McDermott, 1998]. The LTS++ browser-based Integrated Development Environment (IDE) includes many features including an editor with syntax highlighting and static error checking, as well as integrated tools for interactive model testing and debugging, generating and visualizing multiple hypotheses for user-provided observations. The system then translates the LTS++ specification into an AI Planning problem automatically.

The second is through what is called a Mind Map. A Mind Map is a graphical representation of the concepts and relations. The domain knowledge can be encoded by one or more Mind Maps connected by the same concept used in multiple Mind Maps. The Mind Maps can be created in a tool such as FreeMind that produces an XML representation of the Mind Maps and be provided to a system. The system then translates the Mind Maps into an AI Planning problem automatically. It is also possible to learn the causal relation between the concepts in order to build the Mind Maps automatically from scratch or augment or validate existing ones [Hassanzadeh *et al.*, 2019].

4 Core Planning Research

In applying AI Planning techniques we often encounter novel research problems. The following is a list of the contributions to planning research that I was involved with:

Hierarchical Task Network (HTN) planning with preferences [Sohrabi *et al.*, 2009], multi-agent plan recognition [Shvo *et al.*, 2018], top-k planning [Sohrabi *et al.*, 2016b; Katz *et al.*, 2018b], diverse planning [Katz and Sohrabi, 2019], top-quality planning [Katz *et al.*, 2019], and planning and learning [Manikonda *et al.*, 2017; Katz *et al.*, 2018a; Sievers *et al.*, 2019].

5 Planning-based Tools

The final challenge is to address everything else that has to do with developing an AI Planning-based end-to-end system. This includes interaction with the users, transformation of data into observations, automatic translation of the domain knowledge into an AI Planning problem, post-processing the plans, and visualization of the solution. I was involved in developing several tools including LTS++, a tool for the hypothesis generation problem in the large-scale data analysis application [Sohrabi *et al.*, 2016c], as well as the IBM Research Scenario Planning Advisor, a tool for the enterprise risk management application [Sohrabi *et al.*, 2018a].

References

- [Hassanzadeh *et al.*, 2019] Oktie Hassanzadeh, Debarun Bhattacharjya, Mark Feblowitz, Kavitha Srinivas, Michael Perrone, Shirin Sohrabi, and Michael Katz. Answering binary causal questions through large-scale text mining: An evaluation using cause-effect pairs from human experts. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- [Katz and Sohrabi, 2019] Michael Katz and Shirin Sohrabi. Reshaping diverse planning: Let there be light! In *Proceedings of the Heuristics and Search for Domain-independent Planning (HSDIP) workshop at ICAPS*, 2019.
- [Katz *et al.*, 2018a] Michael Katz, Shirin Sohrabi, Horst Samulowitz, and Silvan Sievers. Delfi: Online planner selection for cost-optimal planning. In *Ninth International Planning Competition (IPC-9)*, pages 55–62, 2018.
- [Katz *et al.*, 2018b] Michael Katz, Shirin Sohrabi, Octavian Udrea, and Dominik Winterer. A novel iterative approach to Top-k planning. In *Proceedings of the 28th International Conference on Automated Planning and Scheduling*, 2018.
- [Katz *et al.*, 2019] Michael Katz, Shirin Sohrabi, and Octavian Udrea. Top-quality: Finding practically useful sets of best plans. In *Proceedings of the Heuristics and Search for Domain-independent Planning (HSDIP) workshop at ICAPS*, 2019.
- [Liaskos *et al.*, 2010] Sotirios Liaskos, Sheila A. McIlraith, Shirin Sohrabi, and John Mylopoulos. Integrating preferences into goal models for requirements engineering. In *Proceedings of the 10th International Requirements Engineering Conference (RE)*, pages 135–144, 2010.
- [Liaskos *et al.*, 2011] Sotirios Liaskos, Sheila A. McIlraith, Shirin Sohrabi, and John Mylopoulos. Representing and reasoning about preferences in requirements engineering. *Requirements Engineering*, 16:227–249, 2011. 10.1007/s00766-011-0129-9.

- [Manikonda *et al.*, 2017] Lydia Manikonda, Shirin Sohrabi, Kartik Talamadupula, Biplav Srivastava, and Subbarao Kambhampati. Extracting incomplete planning action models from unstructured social media data to support decision making. In *Proceedings of the Automated Planning and Scheduling (ICAPS) Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*, 2017.
- [McDermott, 1998] Drew V. McDermott. PDDL — The Planning Domain Definition Language. Technical Report TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.
- [Ramírez and Geffner, 2009] Miquel Ramírez and Hector Geffner. Plan recognition as planning. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1778–1783, 2009.
- [Ramírez and Geffner, 2010] Miquel Ramírez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 1121–1126, 2010.
- [Riabov *et al.*, 2015] Anton V. Riabov, Shirin Sohrabi, Daby M. Sow, Deepak S. Turaga, Octavian Udrea, and Long H. Vu. Planning-based reasoning for automated large-scale data analysis. In *Proceedings of the 25th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 282–290, 2015.
- [Shvo *et al.*, 2018] Maayan Shvo, Shirin Sohrabi, and Sheila A. McIlraith. An ai planning-based approach to the multi-agent plan recognition problem. In *Proceedings of the 31st Canadian Conference on Artificial Intelligence (CCAI)*, pages 253–258, 2018.
- [Sievers *et al.*, 2019] Silvan Sievers, Michael Katz, Shirin Sohrabi, Horst Samulowitz, and Patrick Ferber. Deep learning for cost-optimal planning: Task-dependent planner selection. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI-19)*, 2019.
- [Sohrabi and McIlraith, 2009] Shirin Sohrabi and Sheila A. McIlraith. Optimizing Web service composition while enforcing regulations. In *Proceedings of the 8th International Semantic Web Conference (ISWC)*, pages 601–617, 2009.
- [Sohrabi and McIlraith, 2010] Shirin Sohrabi and Sheila A. McIlraith. Preference-based Web service composition: A middle ground between execution and search. In *Proceedings of the 9th International Semantic Web Conference (ISWC)*, pages 713, 729, 2010.
- [Sohrabi *et al.*, 2009] Shirin Sohrabi, Jorge A. Baier, and Sheila A. McIlraith. HTN planning with preferences. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1790–1797, 2009.
- [Sohrabi *et al.*, 2010] Shirin Sohrabi, Jorge Baier, and Sheila McIlraith. Diagnosis as planning revisited. In *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*, pages 26–36, 2010.
- [Sohrabi *et al.*, 2011] Shirin Sohrabi, Jorge A. Baier, and Sheila A. McIlraith. Preferred explanations: Theory and generation via planning. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, pages 261–267, 2011.
- [Sohrabi *et al.*, 2013a] Shirin Sohrabi, Octavian Udrea, and Anton Riabov. HTN planning for the composition of stream processing applications. In *Proceedings of the 23rd International Conference on Automated Planning and Scheduling - Novel Applications Track (ICAPS-13)*, pages 443–451, 2013.
- [Sohrabi *et al.*, 2013b] Shirin Sohrabi, Octavian Udrea, and Anton Riabov. Hypothesis exploration for malware detection using planning. In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI)*, pages 883–889, 2013.
- [Sohrabi *et al.*, 2016a] Shirin Sohrabi, Anton Riabov, and Octavian Udrea. Plan recognition as planning revisited. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3258–3264, 2016.
- [Sohrabi *et al.*, 2016b] Shirin Sohrabi, Anton Riabov, Octavian Udrea, and Otkie Hassanzadeh. Finding diverse high-quality plans for hypothesis generation. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 1581–1582, 2016.
- [Sohrabi *et al.*, 2016c] Shirin Sohrabi, Octavian Udrea, Anton Riabov, and Otkie Hassanzadeh. Interactive planning-based hypothesis generation with LTS++. In *Proceedings of Application Showcase Program at the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [Sohrabi *et al.*, 2017] Shirin Sohrabi, Anton Riabov, and Octavian Udrea. State projection via AI planning. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-17)*, pages 4611–4617, 2017.
- [Sohrabi *et al.*, 2018a] Shirin Sohrabi, Michael Katz, Otkie Hassanzadeh, Octavian Udrea, and Mark D. Feblowitz. Ibm scenario planning advisor: Plan recognition as AI planning in practice. In *Proceedings of Demonstration Track at the the 27th International Joint Conference on Artificial Intelligence (IJCAI-18)*, 2018.
- [Sohrabi *et al.*, 2018b] Shirin Sohrabi, Anton V. Riabov, Michael Katz, and Octavian Udrea. An AI planning solution to scenario generation for enterprise risk management. In *Proceedings of the 32nd National Conference on Artificial Intelligence (AAAI)*, pages 160–167, 2018.
- [Sohrabi *et al.*, 2019] Shirin Sohrabi, Michael Katz, Otkie Hassanzadeh, Octavian Udrea, Mark D. Feblowitz, and Anton Riabov. IBM scenario planning advisor: Plan recognition as AI planning in practice. *AI Commun.*, 32(1):1–13, 2019.
- [Sohrabi, 2012] Shirin Sohrabi. *Customizing the Composition of Web Services and Beyond*. PhD thesis, University of Toronto, 2012.