Towards Automated Planning for Enterprise Services: Opportunities and Challenges

Maja Vukovic[1], Scott Gerard[1], Rick Hull[1], Michael Katz[1], Laura Shwartz[1], Shirin Sohrabi[1], Christian Muise[1], John Rofrano[1], Anup Kalia[1], Jinho Hwang[1], Dang Yabin[2], Ma Jie[2], Jiang Zhuoxuan[2]

[1] IBM T.J. Watson Research Center
{maja, sgerard, lshwartz, ssohrab, rofrano, jinho}@us.ibm.com
{michael.katz1,christian.muise, anup.kalia}@ibm.com
[2] IBM China Research Lab
{dangyb,bjmajie,jzxjiang}@cn.ibm.com

Abstract. Existing Artificial Intelligence (AI) driven automation solutions in enterprises employ machine learning, natural language processing, and chatbots. There is an opportunity for AI Planning to be applied, which offers reasoning about action trajectories to help build automation blueprints. AI Planning is a problem-solving technique, where knowledge about available actions and their consequences is used to identify a sequence of actions, which, when applied in a given initial state, satisfy a desired goal. AI Planning has successfully been applied in a number of domains ranging from space applications, logistics and transportation, manufacturing, robotics, scheduling, e-learning, enterprise risk management, and service composition. In this paper, we discuss experience in building automation solutions that employ AI planning for use in enterprise IT and business services, such as change and event management, migration and transformation and RPA composition. We discuss challenges in adoption of AI planning across the enterprise from implementation and deployment perspectives.

Keywords: Enterprise Services · Change and Event Management · Migration · AI planning

1 Introduction

For the past several years there has been a focus on AI driven automation in services business, with successful solutions that employ Machine Learning (ML), Natural Language Processing (NLP), and chatbots. We believe there is also an opportunity to apply AI Planning in multiple domains e.g., to accelerate migration processes and IT service management (ITSM), speed up creation of RPA solutions and to further enhance chatbots.

AI Planning offers a capability to reason about possible action trajectories to help build automation blueprints. We discuss use cases in services business, proposed framework and challenges to adoption of AI planning.

2 Use Cases

AI Planning offers automated means to schedule, re-plan when needed and manage the design phase of migration processes [8], and can be further used to dynamically assemble (and reassemble) sequences of actions that drive the migration execution process. We also see an opportunity for AI planning in IT Service Management, such as event management, change management, and service management. Similarly, one catalyst for exploring the use of AI Planning in the process specification area is the field of Robotic Process Automation (RPA), and more broadly workflow management.

The general idea of using AI Planning to orchestrate services was explored in the Web Services area [10]. Hoffmann et al. [6] provides a first exploration of how planning can be applied in connection with SAP-based business processes. In task-oriented conversation scenarios, bots are designed to follow a pre-defined dialog flow to respond to user's intents, collect information, take actions and fulfill the user predefined task. But in many business scenarios, tasks are complex and difficult to define as a logical flow, and this opens an opportunity for AI planning.

3 Framework for AI Planning in Enterprise

There are two phases in the application of AI Planning to business domains. At design time: the domain is explored, the problem is specified, stakeholders validate the expected value, the solution is designed, and the implementation is built. Typically at run time, the implementation is executed, plans are created on-demand, and are recreated as unforeseen events occur plan actions are executed, and the implementation is refined as necessary. The design time vs. run time distinction is somewhat variable. Specifically, if the domain is both stable (unchanging) and well understood, it is possible that plans can be generated once during design time and then used repeatedly throughout run time. In this case, a Subject Matter Experts (SME) might review and refine the plan before it is used for repeated execution. Also, in this approach the plans typically include conditionals, so that a single plan can accommodate variations that will arise as the plan is applied to different real-world cases.

Figure 1 illustrates how a Planning capability can be situated in an overall process automation scenario. Starting from the bottom center, the figure shows a library of re-usable configurable items, including micro-robots, decisioning capabilities, and manual tasks. The engines for the configurable items are shown at the bottom, including the ML-based decisioning components, the Robot Engines (from various vendors), and the rules engine. The configurable items are callable using REST APIs, gRPC APIs, or similar. The meta-data about the configurable items is based on a World Model that supports generic business processes (e.g., supporting notions of process flows and conditionals) and also the application at hand (e.g., including data structures for the different business entity types that progress through the process). There is also meta-data about the client environment (e.g., system configuration information).

When applying AI Planning for process specification there is a design-time variant, where the plan is created in advance and includes conditionals, and a run-time variant, where planning is preformed dynamically, at the beginning and/or in the middle of process execution. The design-time variant is appropriate when developing a process that will be run many times (e.g., invoice processing, supply chain management, report generation) or where a manual validation and refinement of the automatically generated plan is desired. The run-time variant is useful for handling exceptional situations. Figure 1 shows Planning Engines in two places, corresponding to the use of design-time or run-time planning.

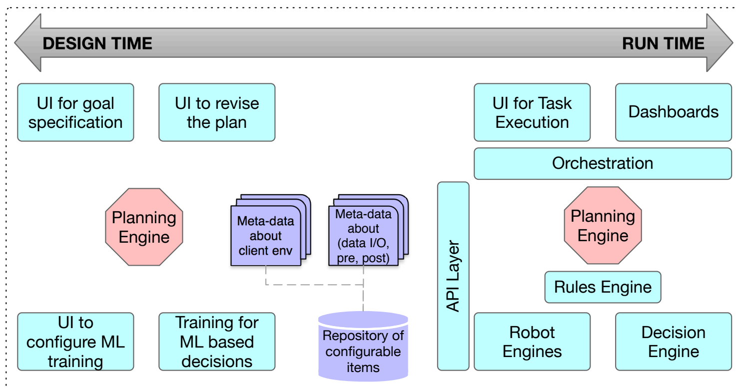

Fig. 1. High-level framework

Finally, the top layer of the figure shows the main user interface capabilities of the framework. These include, for design-time planning, a UI for specifying the goals to be achieved, and also for reviewing and refining the generated plans and process schemas. For both design-time and run-time planning there are UIs for end-user task performance and for monitoring and reporting.

4 Challenges in AI Planning for Enterprise

Non-experts rarely are familiar with the existing planning tools and their formalisms. For example, tools that provide a provably optimal solution might be of interest for some tasks, while tools that provide a solution of increased quality in an anytime manner can be desired in other cases. Thus, there is a clear benefit in removing the burden of choosing the right tool from the non-expert user.

The work on planning portfolios, where multiple planners for the same computational problem are exploited to derive a meta-planner for that computational problem is one step in that direction [4, 9, 12]. Across computational problems, aside of educating the user, not much has been done so far. One possible step in that direction is creating a single entry point that receives a planning task and returns solutions for multiple computational problems, emphasizing the differences in solution quality, consumed resources, and solution guarantees. That way, the user will be able to make an intelligent decision about the actual computational problem and the respective (meta-)planner to use.

Per computational problem, the performance issues may be tackled by introducing additional planners, and making the choice of the actual planner task dependent. Performance can improve gradually, with every additional planner added to the meta-planner.

Since planning domains often ignore large parts of the actual problem at hand, all too often, the obtained solutions are not fully applicable in the real world. Further, since planning tools are not easily maintained by generalists, it might be beneficial to minimize the actual planning at the critical path of an application.

To ease the entry into planning for the general crowd, it might be beneficial to focus on a small set of domains of specific interest, such as IT service management. In these domains, specific tools can be created to extract the knowledge needed in order to create the planning model. Creating such tools will reduce the barrier to using planners for these applications.

Difficulty of modeling planning tasks is one of the major obstacles to using planners, there are research efforts on developing tools that help with generic, domain-independent modeling, such as itSIMPLE [14]. Other possibilities may include the use of Mind Maps, a graphical representation of concepts and relations, to address the knowledge engineering challenge. The domain knowledge can be encoded by one or more Mind Maps connected by the same concept used in multiple Mind Maps. The system can then translates the Mind Maps into an AI planning problem automatically [13]. 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 [5].

Another area for research is how to develop systems that automatically learn the actions that comprise the domain descriptions. The goal would be to learn both the actions themselves coupled with preconditions and effects. Predicates would need to be consistent across all the atomic actions. Some early work is being done on learning the actions from SMEs (e.g. generating high level descriptions and approximated preconditions and effects, prompting the SME to refine them); business process information is extracted from text-based descriptions using natural language processing [11, 7, 3].

5 Conclusion

We presented use cases of AI planning as means of automating IT and Business services. We introduced key elements of an overarching framework, which brings together planning at design and runtime phases. We then outline key research challenges to full adoption of AI planning for automation in the Enterprise, including lack of tooling that enables seamless and automated development of world models, coupled with generic planning models. Future work will focus on integration of planning and learning.

Acknowledgments We thank our colleagues: Stefan Pappe, Arvind Viswanathan, Valentina Salapura, Sridrar Thiruvengadam, Boby Philip, Sharon Alvarado Brenes, Joaquin Eduardo Bonilla Arias, and Sussana Ting.

References

1. Ninth International Planning Competition (IPC-9): planner abstracts (2018)
2. Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019). AAAI Press (2019)
3. van der Aa, H., Leopold, H., Reijers, H.A.: Detecting inconsistencies between process models and textual descriptions. In: Business Process Management - 13th International Conference, BPM 2015, Innsbruck, Austria, August 31 - September 3, 2015, Proceedings. pp. 90–105 (2015)
4. Cenamor, I., de la Rosa, T., Fernández, F.: IBaCoP-2018 and IBaCoP2-2018. In: Ninth International Planning Competition (IPC-9): planner abstracts [1], pp. 9–10
5. Hassanzadeh, O., Bhattacharjya, D., Feblowitz, M., Srinivas, K., Perrone, M., Sohrabi, S., Katz, M.: Answering binary causal questions through large-scale text mining: An evaluation using cause-effect pairs from human experts. In: IJCAI19 (2019)
6. Hoffmann, J., Weber, I., Kraft, F.M.: Sap speaks pddl: Exploiting a software-engineering model for planning in business process management. J. Artif. Int. Res. **44**(1), 587–632 (2012), http://dl.acm.org/citation.cfm?id=2387933.2387946
7. Hull, R., Nezhad, H.R.M.: Rethinking BPM in a cognitive world: Transforming how we learn and perform business processes. In: BPM 2016, Rio de Janeiro, Brazil, September 18-22, 2016. Proceedings (2016)
8. Jackson, M., Rofrano, J.J., Hwang, J., Vukovic, M.: Blueplan: A service for automated migration plan construction using AI. In: Service-Oriented Computing - ICSOC 2018 Workshops, Hangzhou, China, November, 2018, Revised Selected Papers (2018)
9. Katz, M., Sohrabi, S., Samulowitz, H., Sievers, S.: Delfi: Online planner selection for cost-optimal planning. In: Ninth International Planning Competition (IPC-9): planner abstracts [1], pp. 57–64
10. Narayanan, S., McIlraith, S.A.: Simulation, verification and automated composition of web services. In: Proceedings of the 11th International Conference on World Wide Web. pp. 77–88. WWW '02, ACM, New York, NY, USA (2002)
11. Nezhad, H.R.M., Akkiraju, R.: Towards cognitive BPM as the next generation BPM platform for analytics-driven business processes. In: Business Process Management Workshops - BPM 2014 International Workshops, Eindhoven, The Netherlands, September 7-8, 2014, Revised Papers. pp. 158–164 (2014)
12. Sievers, S., Katz, M., Sohrabi, S., Samulowitz, H., Ferber, P.: Deep learning for cost-optimal planning: Task-dependent planner selection. In: Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019) [2]
13. Sohrabi, S., Katz, M., Hassanzadeh, O., Udrea, O., Feblowitz, M.D., Riabov, A.: IBM scenario planning advisor: Plan recognition as AI planning in practice. AI Commun. **32**(1), 1–13 (2019)
14. Vaquero, T.S., Romero, V., Tonidandel, F., Silva, J.R.: itSIMPLE 2.0: An integrated tool for designing planning domains. In: Boddy, M., Fox, M., Thiébaux, S. (eds.) Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling (ICAPS 2007). pp. 336–343. AAAI Press (2007)