Process Mining-Based Goal Recognition (Extended Abstract)*

Zihang Su[1,*,†]

[1] The University of Melbourne

Abstract

The process mining-based goal recognition system can infer the intention of an autonomous agent based on its historical behaviors. The developed system can learn process models using process discovery techniques and then use conformance diagnostics between the constructed process models and a new observation to formulate a probability distribution over a range of possible goals the agent attempts to achieve. Compared to other state-of-the-art goal recognition approaches, the proposed approach does not rely on handcrafted domain knowledge and, thus, is applicable in many real-world scenarios. Combined with re-learn mechanism and process model repair techniques, the developed system can continuously solve goal recognition problems and is robust to environmental changes.

Keywords

goal recognition, process mining, environmental change, process model repair

1. Introduction

Goal Recognition (GR) techniques aim to infer the intentions of an autonomous agent according to the observed actions of that agent [1]. GR techniques play an important role in many areas, such as the support of adversarial reasoning [2], trajectory prediction [3], and human-computer collaboration [4].

Three concepts are inherent to the understanding of GR: A plan is a sequence of actions that were or should be taken to achieve a goal; An agent, such as a robot or human, follows plans to accomplish goals; A GR system is a software that implements a GR technique capable of inferring the goals of agents based on partial knowledge about the plans (partially observed plans). When a GR system analyzes actions executed by an agent, it aims to forecast the full plan that the agent is following and, hence, the goal that will be achieved after completing the plan.

The existing GR techniques can be classified into two main categories: 1) Observations of an agent's actions are "matched" to a plan (the one judged to being carried out by the agent) in a pre-defined library encoding the standard operational procedures of the domain [1], namely, plan library-based approaches; 2) Appealing to the principle of rational behavior, an agent is assumed to be taking the "optimal" path to the goal: the more rational a behavior appears

ICPM 2022 Doctoral Consortium and Tool Demonstration Track
*Corresponding author.
zihangs@student.unimelb.edu.au (Z. Su)
0000-0003-1039-7462 (Z. Su)

towards a goal, the more likely such goal is the agent's goal. Ramirez and Geffner [5] have sparked a plethora of approaches not needing any a priori set of plans. These approaches perform GR by exploiting planning techniques to automatically generate plans relative to a domain theory, namely, planning-based approaches.

The challenge for plan library-based approaches is in obtaining a set of standard plans and hand-coding an informative model (plan library) that can represent the standard plans. Besides, these approaches do not accommodate uncertainty (cannot generalize to the observations that are not pre-stored in the plan library). For planning-based approaches, even though specifying domain models could be less effort demanding than hand-coding plan libraries, and "new" plans can be found, acquiring domain models is far from trivial due to the difficulty of defining models using standard declarative languages [6]. It is especially challenging to acquire domain models of real-world environments. Hence, planning-based approaches are difficult to apply in real-world scenarios.

2. Proposed Solutions

RQ1 is solved and verified with a published paper [7], explained in 2.1, while RQ2 is in progress [8], and two possible solutions are explained in 2.2

2.1. Answer to RQ1

We proposed a GR approach based on process mining techniques, namely, the process mining (PM-)based approach, to learn knowledge (models) by observing agents' behaviors, then the PM-based GR approach can recognize agents' goals according to newly observed behaviors and the learned knowledge [7]. In this approach, firstly, we collect the previous sequences of actions executed by the agent, and the action sequences for achieving the same goals are grouped and converted to event logs; thus, multiple event logs are obtained, and each event log records the plans for achieving a particular goal. Secondly, we apply the process discovery technique [9, 10] to mine Petri nets (to represent the knowledge models) from the obtained event logs. Thirdly, we apply conformance checking techniques [11] to measure the costs of alignments between the newly observed action sequence and the learned models (Petri nets). Finally, we construct a probability distribution over all goal candidates according to the costs of alignments.

2.2. Answer to RQ2

One of the possible solutions for RQ2 is to re-learn the knowledge model (re-mine Petri nets) based on the behaviors of the agent acting in the new environment. As we assume the full model of the underlying environment is unknown, the environmental change can not be directly observed. Therefore, we need to construct the mechanism for deleting the environmental change and deciding when to re-learn from which observed dataset. The recognition accuracy score can be an indicator of environmental changes (if the underlying environment changes, the agent may act differently in the new environment, which causes the accuracy score to drop). The dataset used for re-learning knowledge models is the feedback from each "single-shot" GR task. Another possible solution is to apply the process model repair techniques [12, 10] for adjusting the knowledge models (Petri-nets) when the GR system receives negative feedback (the inferred goal is not the ground truth). A challenge for applying the process model repair technique is to tackle the problem of deleting a behavior from a knowledge model.

3. Evaluation Methodology

We compared the PM-based GR approach with other state-of-the-art GR approaches [5, 13] using the planning domains[1] as dataset. The result shows that, compared to other approaches, the PM-based GR approach has a similar GR performance in terms of recall and a faster response time. We conducted GR experiments on three real-world domains which obtained from three BPI challenge logs, namely activities of daily living,[2] building permit applications,[3] and environmental permit applications.[4] The results verified that the PM-based GR approach is applicable in real-world scenarios; for details refer to [7]. For RQ2, since continuous GR over changing environments is a new problem, commonly used experimental datasets and benchmarks are not available. Thus, we built a simulator [8] that can generate experimental data based on classical planning domains.[1] We will conduct experiments on the generated datasets to evaluate the GR performance before and after the environmental changes. We will compare different re-learn mechanisms and model repair techniques to study if there exists an optimal solution for all scenarios.

4. Conclusion

This research project utilizes the existing process mining-related techniques [9, 10, 12, 11] to implement a GR approach that can automatically learn the knowledge model, and overcomes the limitations of for the state-of-the-art GR techniques [1, 5, 13]. Additionally, the research aims to provide a possible solution to the problem of GR in non-stationary environments [14, 15].

Acknowledgments

Thanks to my supervisors, Dr. Artem Polyvyanyy, Dr. Nir Lipovetzky, Prof. Sebastian Sardina, and Dr. Nick van Beest. Thanks to the funding institutions, the University of Melbourne, Chinese Scholarship Council, and Data61 CSIRO.

References

[1] H. A. Kautz, J. F. Allen, Generalized plan recognition, in: AAAI, Morgan Kaufmann, 1986, pp. 32–37.

[2] G. Sukthankar, K. P. Sycara, A cost minimization approach to human behavior recognition, in: AAMAS, ACM, 2005, pp. 1067–1074.

[1] https://github.com/pucrs-automated-planning/goal-plan-recognition-dataset/

[2] https://doi.org/10.4121/uuid:01eaba9f-d3ed-4e04-9945-b8b302764176

[3] https://doi.org/10.4121/uuid:31a308ef-c844-48da-948c-305d167a0ec1

[4] https://doi.org/10.4121/uuid:26aba40d-8b2d-435b-b5af-6d4bfbd7a270

[3] J. Firl, Q. Tran, Probabilistic maneuver prediction in traffic scenarios, in: ECMR, Learning Systems Lab, AASS, Örebro University, 2011, pp. 89–94.

[4] N. Lesh, C. Rich, C. L. Sidner, Using plan recognition in human-computer collaboration, in: UM99 User Modeling, Springer, 1999, pp. 23–32.

[5] M. Ramírez, H. Geffner, Probabilistic plan recognition using off-the-shelf classical planners, in: AAAI, AAAI Press, 2010.

[6] P. Haslum, N. Lipovetzky, D. Magazzeni, C. Muise, An Introduction to the Planning Domain Definition Language, Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers, 2019.

[7] A. Polyvyanyy, Z. Su, N. Lipovetzky, S. Sardiña, Goal recognition using off-the-shelf process mining techniques, in: AAMAS, International Foundation for Autonomous Agents and Multiagent Systems, 2020, pp. 1072–1080.

[8] Z. Su, A. Polyvyanyy, N. Lipovetzky, S. Sardina, N. van Beest, Grace: A simulator for continuous goal recognition over changing environments, in: PMAI-IJCAI, 2022, in press.

[9] A. Augusto, R. Conforti, M. Dumas, M. L. Rosa, A. Polyvyanyy, Split miner: automated discovery of accurate and simple business process models from event logs, Knowl. Inf. Syst. 59 (2019) 251–284.

[10] W. M. P. van der Aalst, V. A. Rubin, H. M. W. Verbeek, B. F. van Dongen, E. Kindler, C. W. Günther, Process mining: a two-step approach to balance between underfitting and overfitting, Softw. Syst. Model. 9 (2010) 87–111.

[11] W. M. P. van der Aalst, A. Adriansyah, B. F. van Dongen, Replaying history on process models for conformance checking and performance analysis, WIREs Data Mining Knowl. Discov. 2 (2012) 182–192.

[12] A. Polyvyanyy, W. M. P. van der Aalst, A. H. M. ter Hofstede, M. T. Wynn, Impact-driven process model repair, ACM Trans. Softw. Eng. Methodol. 25 (2017) 28:1–28:60.

[13] R. F. Pereira, N. Oren, F. Meneguzzi, Landmark-based approaches for goal recognition as planning, Artif. Intell. 279 (2020).

[14] D. Bryce, J. Benton, M. W. Boldt, Maintaining evolving domain models, in: IJCAI, IJCAI/AAAI Press, 2016, pp. 3053–3059.

[15] T. Chakraborti, S. Sreedharan, Y. Zhang, S. Kambhampati, Plan explanations as model reconciliation: Moving beyond explanation as soliloquy, in: IJCAI, ijcai.org, 2017, pp. 156–163.