Learning Planning Domains for Intelligent Transport Systems*

Simona Gugliermo[1]

Abstract— This short paper presents an ongoing industrial PhD project focused on learning domain models in the transportation industry. We seek to develop a two-phase learning process that leverages Behavior Trees to involve human planning experts in the learning loop. In this paper, we outline our research questions, discuss the initial steps we have taken in our research, and highlight the challenges we expect to face moving forward.

I. PROBLEM STATEMENT

Automated planning methods (AI planning) have the potential to introduce disruptive changes in industrial applications. They can automate critical processes such as resource allocation, task scheduling, and contingency management that are currently carried out manually. This need for high-quality plans and effective execution is particularly crucial in industries investing in autonomous transport vehicles, including mining, construction, and material transport.

In Scania's current industrial applications, plans are usually created manually by specialists who possess extensive knowledge about the particular field of application. Even when AI planners are used, they still require manually provided domain-specific information in the form of action models to generate plans. These models encode preconditions and effects of actions in the *planning domain* [1]. Obtaining accurate and complete domain descriptions is a challenging task that requires significant knowledge engineering efforts, hindering the use of AI planning techniques in real-world scenarios. To overcome this challenge, our research focuses on automatically extracting and formalizing human planning expert knowledge from historical executed traces to generate planning domains. Furthermore, the use of learning can facilitate the transfer of knowledge from humans to machines as inferring domain knowledge from simulation and execution is one of the main challenges in this context.

II. PROPOSED APPROACH AND RESEARCH QUESTIONS

The objective of this PhD project is to enhance the current level of manual planning by utilizing machine learning while incorporating human expertise. Currently, the so called *fleet transport managers* are responsible for manually creating plans. We propose a two-step learning approach, shown in Figure 1, that utilizes an intermediate human readable representation. Introducing an intermediary representation has three key benefits. Firstly, it simplifies the learning process by bridging the gap between execution traces and symbolic action representations. Secondly, it leverages human experts knowledge and experience. Lastly, it enables human inspection and validation, thus enhancing transparency in the learning process.

To achieve these benefits, we propose utilizing Behavior Trees (BTs), which, according to Colledanchise and Ögren [2], are known to be comprehensible by humans. Indeed, their hierarchical structure provides a more intuitive and easily understandable representation when compared to planning domains. In detail, the proposed approach involves generating planning domain models by using execution traces from manually computed plans. These traces are fed into a machine learning process that employs data mining and mathematical logic techniques to learn a BT. Once the BT is validated by human experts, it is combined with causality data extracted from observed data to automatically generate the planning domain. This domain, along with a specific planning problem, is then inputted into an AI planner to generate an effective plan. Overall, this project centers on three specific research questions (RQs):

RQ1: *How to extract a BT from historical execution traces?*
This research question addresses the need to develop techniques to learn the implicit knowledge encoded in execution traces in a human-readable way. This is important because it can lead to more accurate and efficient planning processes. While similar to Learning from Demonstation (LfD), learning a correct and concise BT from executed traces remains a significant challenge despite prior research [4].

RQ2: *How to enable humans to validate the BT?*
This research question addresses the need to incorporate human experts into the learning phase to ensure that the learned BT is correct (verification), satisfies the requirements (validation), and is free of errors and bugs (debugging). The goal is to aid human experts in evaluating and improving BTs in a more efficient and immediate manner.

RQ3: *How to obtain a domain theory from the BT?*
This research question addresses the need to automate the generation of planning domain theories from machine learning-generated intermediate representations. This involves identifying the preconditions necessary for an action to be performed and the effects of an action on a state variable. By automating the process of extracting the planning domain, we expect to significantly reduce the workload of human experts who are responsible for manually creating planning domain models for real-world applications.

[1]Author is at the Center for Applied Autonomous Sensor Systems (AASS), Örebro University, S-701 82 Örebro, Sweden and at Intelligent Transport Systems, Scania CV AB, S-151 87 Södertälje, Sweden
simona.gugliermo@oru.se

Fig. 1. Visual representation of the proposed approach for addressing the three fundamental Research Questions (RQs). By mapping the approach to these questions, the figure highlights the key steps involved in the proposed method and how they align with the research objectives thus providing a clear and concise overview.

III. CURRENT RESULTS

Our primary research efforts have been focused on addressing RQ1. We have successfully developed a framework for extracting and representing implicit action knowledge from execution traces [5]. Our method employs decision trees (DTs) and mathematical logic techniques to learn a BT, which accurately captures the logical elements implicit in the historical data and generates correct behavior in a simulated domain.

In relation to RQ2, although interactive tools are already available for validating behavior trees, we are currently investigating the possibility of using quantifiable metrics to formally evaluate a BT's properties, thus aiding the human expert in a more efficient and immediate evaluation.

To address RQ3, it is essential to infer causal relationships between actions and state variables, which are encoded in the form of preconditions and effects. We propose using the DT obtained from our framework to infer preconditions, as it expresses the conditions necessary for an action to be performed. However, derive the effects of an action on a state is a more complex task. To address this, we are currently investigating an approach that utilizes observable historical data inspired by the Difference-in-Differences (DID) technique [3].

IV. NEXT STEPS

Based on our current research progress, there are several areas that require further investigation and development. Regarding RQ2, we will continue exploring the definition and use of formal metrics for evaluating BTs. We aim to identify the most useful and effective metrics for assessing a BT's properties and plan to conduct empirical studies to validate their usefulness. Additionally, we will explore ways

to incorporate these metrics into our existing framework to make it more effective and efficient.

For RQ3, we will continue to investigate the use of observable historical data to infer the effects of actions on state variables. Specifically, we will further develop our approach inspired by the DID technique, aiming to identify and estimate the causal effects of actions on state variables and integrate them into our framework.

Our ultimate goal is to advance the state of the art in learning domain models for AI planning methods and develop practical tools and techniques that can aid manual planning experts in their work. We aim to bridge the gap between human expertise and machine learning, ultimately improving the efficiency and accuracy of planning processes in real-world applications.

V. ACKNOWLEDGMENT

REFERENCES

[1] M. Ghallab, D. Nau, and P. Traverso. "Automated Planning and Acting". Cambridge University Press, Cambridge, 2016. ISBN 978-1-107-03727-4.

[2] M. Colledanchise and P. Ögren. "Behavior Trees in Robotics and AI". CRC Press, 2018.

[3] J.A. Hausman. "Specification Tests in Econometrics". Econometrica, vol. 46, no. 6, pp. 1251-1271, 1978. Publisher: Wiley, Econometric Society. ISSN: 00129682, 14680262.

[4] A. Wathieu, T.R. Groechel, H.J. Lee, C. Kuo, and M.J. Mataric. "RE:BT-Espresso: Improving Interpretability and Expressivity of Behavior Trees Learned from Robot Demonstrations". In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2022.

[5] S. Gugliermo, E. Schaffernicht, C. Koniaris, and F. Pecora, "Learning Behavior Trees From Planning Experts Using Decision Tree and Logic Factorization". In IEEE Robotics and Automation Letters, vol. 8, no. 6, pp. 3534-3541, 2023, doi: 10.1109/LRA.2023.3268598.