

Chapter 29

Improving Automated Planning with Machine Learning

Susana Fernández Arregui

Universidad Carlos III de Madrid, Spain

Sergio Jiménez Celorrio

Universidad Carlos III de Madrid, Spain

Tomás de la Rosa Turbides

Universidad Carlos III de Madrid, Spain

ABSTRACT

This chapter reports the last machine learning techniques for the assistance of automated planning. Recent discoveries in automated planning have opened the scope of planners, from toy problems to real-world applications, making new challenges come into focus. The planning community believes that machine learning can assist to address these new challenges. The chapter collects the last machine learning techniques for assisting automated planners classified in: techniques for the improvement of the planning search processes and techniques for the automatic definition of planning action models. For each technique, the chapter provides an in-depth analysis of their domain, advantages and disadvantages. Finally, the chapter draws the outline of the new promising avenues for research in learning for planning systems.

INTRODUCTION

Automated Planning (AP) is the branch of Artificial Intelligence (AI) that studies the computational synthesis of sets of actions to carry out a given task (Ghallab et al., 2004). AP appeared in the late '50s from converging studies into state-space search, theorem proving and control theory to solve the practical needs of robotics. The Stanford Institute Problem Solver STRIPS (Fikes & Nilsson, 1971), developed to be the planning component for controlling the autonomous robot Shakey (Nilsson, 1984), perfectly illustrates the interaction of these influences. Since the Shakey's days up to now, AP has created efficient planners and well accepted standards for representing and solving the planning tasks. In the last

DOI: 10.4018/978-1-60566-766-9.ch029

years, AP systems have been successfully applied for controlling underwater vehicles (McGann et al., 2008), for the management of fire extinctions (Castillo et al., 2006), for the planning of space missions (Bresina et al., 2005), etc . But despite of all these successful examples, the application of AP systems to real-world problems is still complicated, mainly due to the following two factors. First, the search for a solution plan in AP is a PSpace-complete¹ problem (Bylander, 1991). One can easily find planning problems that overwhelm the capacities of the off-the-shelf planners. Second, AI planners need to be fed with accurate description of the planning tasks. These descriptions consist of a model of the actions that can be carried out in the environment together with a specification of the state of the environment and the goals to achieve. Knowing in advance this information is unfeasible in most of the real-world problems.

The following sections describe different approaches for applying Machine Learning (ML) to assist AP in order to overcome these two problems. These approaches learn either heuristics (also called control knowledge) for handling the search complexity or action models. The chapter is organized as follows. The first section explains the basic concepts of AP. Second section describes the common issues about using ML for AP. Third section reviews the last techniques for automatically learning AP control knowledge. Fourth section reviews the last advances for the learning of AP action models. The fifth section depicts the current challenges of ML for AP. And finally, section sixth discusses some conclusions.

BACKGROUND

An AP task is defined by two elements: (1) a set of actions that represents the state-transition function of the world (*the planning domain*) and (2), a set of facts that represent the initial state together with the goals of the AP task (*the planning problem*). These two elements are typically represented in languages coming from the first-order logic. In the early days of AP, STRIPS was the most popular representation language. In 1998 the Planning Domain Definition Language (PDDL) was developed for the first International Planning Competition (IPC). Since that date, PDDL has become the standard representation language for the AP community. According to the PDDL specification (Fox & Long, 2003), an action in the planning domain is represented by: (1) the action preconditions, a list of predicates indicating the facts that must be true so the action becomes applicable and (2) the action effects, which is a list of predicates indicating the changes in the state after the action application. Like STRIPS, PDDL follows the closed world assumption to solve the frame problem. Regarding this assumption, what is not currently known to be true, is false.

Before the mid '90s, automated planners could only synthesize plans of no more than 10 actions in an acceptable amount of time. During those years, planners strongly depended on speedup techniques, such as learning control knowledge, for solving interesting AP problems. In the late 90's, a significant scale-up in planning took place due to the appearance of the reachability planning graph (Blum & Furst, 1995) and the development of powerful domain independent heuristics (Hoffman & Nebel, 2001; Bonet & Geffner, 2001). Planners using these advances can often synthesize 100-action plans just in seconds. So, at the present time, there is not such dependence on ML for solving AP problems. However, there is a renewed interest in learning for AP motivated by two factors: (1) IPC-2000 showed that knowledge-based planners scale up dramatically better than domain-independent planners. The development of ML techniques that automatically define the kind of knowledge that humans put in these planners would bring great advances to the field. With this aim, the IPC-2008 introduced a specific track for learning-

20 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/improving-automated-planning-machine-learning/37006

Related Content

Empirical Studies on the Functional Complexity of Software in Large-Scale Software Systems

Yingxu Wang and Vincent Chiew (2011). *International Journal of Software Science and Computational Intelligence* (pp. 23-42).

www.irma-international.org/article/empirical-studies-functional-complexity-software/60747

Rough and Soft Set Approaches for Attributes Selection of Traditional Malay Musical Instrument Sounds Classification

Norhalina Senan, Rosziati Ibrahim, Nazri Mohd Nawi, Iwan Tri Riyadi Yanto and Tutut Herawan (2012). *International Journal of Software Science and Computational Intelligence* (pp. 14-40).

www.irma-international.org/article/rough-soft-set-approaches-attributes/72878

The Formal Design Model of a Real-Time Operating System (RTOS+): Conceptual and Architectural Frameworks

Yingxu Wang, Cyprian F. Ngolah, Guangping Zeng, Philip C.Y. Sheu, C. Philip Choy and Yousheng Tian (2010). *International Journal of Software Science and Computational Intelligence* (pp. 105-122).

www.irma-international.org/article/formal-design-model-real-time/43900

Cognitive Location-Aware Information Retrieval by Agent-Based Semantic Matching

Eddie C. L. Chan, George Baciu and S. C. Mak (2012). *Breakthroughs in Software Science and Computational Intelligence* (pp. 335-345).

www.irma-international.org/chapter/cognitive-location-aware-information-retrieval/64616

Design of Low-Power High-Speed 8 Bit CMOS Current Steering DAC for AI Applications

Banoth Krishna, Sandeep Singh Gill and Amod Kumar (2022). *International Journal of Software Science and Computational Intelligence* (pp. 1-18).

www.irma-international.org/article/design-of-low-power-high-speed-8-bit-cmos-current-steering-dac-for-ai-applications/304801