AI Education: Adaptive Planning

Joshua Eckroth (Stetson University; jeckroth@stetson.edu)
DOI: 10.1145/3175502.3175508

Introduction

In this column, we focus on designing assignments and projects that make use of planning engines. Planning has been one of the pillars of artificial intelligence since the origin of the field, and the research community remains active, as evidenced by competitions such as the IPC and conferences such as ICAPS.

Planning's practical applications and the increasing sophistication of planning engine design warrant its place in the classroom. However, planning's historical roots bring some negative side effects: planning is not "hot" like deep learning, thus reducing student interest; and many examples, assignments, and projects related to planning may appear stale and uninteresting since they may well be decades old.

An Adaptive Approach

In my experience, students are more engaged with an assignment or project if it seems "fresh," or stated less colloquially, adaptive to the student, local community, or zeitgeist. The material does not need to be entirely novel or make use of the most recent advances in the field. But to fully engage a student, we must find ways to position a project as tailored to the individual or class, or addressing a societal need, so that they feel there is some wider significance in completing it.

An Example: Git Planner

I have found success in engaging students by asking them to build planners for commonplace but distinctive skills and systems. For example, I designed a project called *Git Planner* (Neller et al., 2017), available online, that modeled the logic of the Git version control system and commands such as *git add*, *git revert*, etc. The planner would find a sequence of Git commands to take a repository's initial state and transform it into a user-provided goal state, e.g., that all files are committed or a certain file was reverted to a former state. My goal was to pique interest by asking students to model and automate the use of a famously difficult tool. The project was fairly narrow in its ambitions, modeling only a few common Git commands, but the students knew they were embarking on a task that had little precedent. After the course, a senior student recently completed an expanded version of the project for his year-long capstone.

A Modern Curriculum

Planning's long history has yielded a broad and mature set of techniques. A deep investigation would require at least a full course. Many schools do not have the opportunity to introduce a course focused on planning. Instead, I recommend covering a few significant topics to introduce students to planning and its applications. The idea of *planning as search* may be covered early in an AI course. Pathfinding for characters in a game is a classic example of simple planning, and can be applied to modern games to ensure student interest. Planning a sequence of actions, such as assembling a desktop computer from scratch, may be accomplished with a partial order planner. This example exercise requires that students investigate exactly how a computer can be assembled, which actions depend on which other actions, and which actions can be done in any order.

Students with a background in probability may be ready to work with Markov decision processes and partially-observable Markov decision processes (Kaelbling, Littman, & Cassandra, 1998), abbreviated POMDP, in which the result of an action is nondeterministic and the agent might not even be completely sure of its present state. Here, actions that allow the agent to learn more about the environment may be equally useful as actions that help achieve long-term goals. Relevant applications are numerous (Cassandra, 1998). An adaptive application of probabilistic plan-

ning could address the contemporary focus on cybersecurity to help design a secure network, including where to introduce a firewall, which services are only accessible by passwords, password complexity requirements to reduce the chance of brute force hacking, and the likelihood of exploiting out-of-date vulnerable software. A hacker's behavior may be simulated with a POMDP planner and the network design and policies may be improved in order to limit the hacker's probability of success.

Resources

Adaptive planning assignments and projects need not be developed from scratch. The Model AI Assignments repository, from years 2010 to 2017, includes seven projects that make use of planning. Planning competitions release challenge datasets and clearly describe the intended outcomes. Competitions are often highly motivating to students and may be considered adaptive, as described above, since they are relevant for a short period of time and students participate alongside other teams across the globe – in other words, it matters that the students are working on this project at this time, and the project requires the use of state-of-the-art techniques. However, planning competitions may be too challenging for introductory projects. Some planning engines such as Fast Downward include example problems in their documentation, which can serve as a starting point for new ideas. Finally, adaptive assignments may be found by considering an apparently disconnected field of study, such as cybersecurity, marketing, or medicine just to name a few, and attempting to solve a simplified form of their everyday concerns with planning technology.

Conclusion

Adaptive planning projects are more likely to keep students engaged because adaptive projects take advantage of a unique time, place, or other context that gives the project of sense of relevance and urgency that is typically lacking from traditional textbook curricula. Planning, like other traditional subfields of artificial intelligence, particularly need adaptive materials to compete with the ever-changing trends in the popular understanding of artificial intelligence and machine learning.

References

Cassandra, A. R. (1998). A survey of POMDP applications. In *Working notes of AAAI 1998 fall symposium on planning with partially observable markov decision processes* (Vol. 1724).

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, *101*(1), 99–134.

Neller, T. W., Eckroth, J., Reddy, S., Ziegler, J., Bindewald, J., Peterson, G., ... Karaman, S. (2017). Model AI assignments 2017. In *Proceedings of the seventh symposium on educational advances in artificial intelligence* (pp. 4822–4824).

Joshua Eckroth is an assistant professor at Stetson University and chief architect of i2k Connect. His research interests lie in abductive reasoning and belief revision as well as computer science pedagogy. Eckroth holds a Ph.D. in computer science from the Ohio State University. He is coeditor of AITopics.