

Adaptive Metareasoning for Bounded Rational Agents

Justin Svegliato and Shlomo Zilberstein

College of Information and Computer Sciences

University of Massachusetts Amherst

{jsvegliato,shlomo}@cs.umass.edu

In computational approaches to bounded rationality, metareasoning enables intelligent agents to optimize their own decision-making process in order to produce effective action in a timely manner. While there have been substantial efforts to develop effective meta-level control for anytime algorithms, existing techniques rely on extensive offline work, imposing several critical assumptions that diminish their effectiveness and limit their practical utility in the real world. In order to eliminate these assumptions, adaptive metareasoning enables intelligent agents to adapt to each individual instance of the problem at hand without the need for significant offline preprocessing. Building on our recent work, we first introduce a model-free approach to meta-level control based on reinforcement learning. We then present a meta-level control technique that uses temporal difference learning. Finally, we show empirically that our approach is effective on a common benchmark in meta-level control.

The full paper is available online. Please use the following link. I am sorry I could not create a PDF/A version and was not able to upload the file.

<http://rbr.cs.umass.edu/shlomo/papers/SZijcaiAEGAP18.pdf>