Reinforcement Learning at the Cognitive Level in a Belief, Desire, Intention UAS Agent

David Sacharny, Thomas C. Henderson,
Michael Cline and Ben Russon
University of Utah

UUCS-20-13

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

23 October 2020

Abstract

We have proposed a lane-based approach to handle large-scale Unmanned Aircraft Systems (UAS) traffic management. Given the number of UAS in the airways simultaneously, it is necessary to imbue each UAS with a learning capability at the cognitive level so that it can optimize its performance in the face of weather, airway congestion and other contingencies. Here we describe a Belief, Desire, Intention (BDI) architecture for the representation and reasoning over cognitive states, where these beliefs include goals for the UAS such as staying in its lane, on heading and at the desired speed. Such goals are represented as logical propositions, and if false, they may be selected as goals (intentions). We apply reinforcement learning to optimize the selection of a plan to achieve the goal of the UAS. That is, a policy is determined which given the cognitive state of the UAS, including for example weather conditions, a plan is selected which achieves the goal with minimal cost and maximal effectiveness.