MIT LINCOLN LABORATORY

The Bulletin

News, events, and notices transmitted weekly

August 28, 2020

Mission-ready Reinforcement Learning

Artificial Intelligence Technology | Lincoln Laboratory

Reinforcement learning is an exciting emergent technology with strong applicability across a range of Department of Defense (DoD) mission areas. It is a form of machine learning where an agent is trained via rewards and punishments based on its performance. Deep reinforcement learning (DRL) has shown the capacity to find effective strategies in complex and high-dimensional environments, such as learning to play strategy games like StarCraft or Defense of the Ancients. As a result, the DoD is investing in simulation environments for training autonomous agents to solve challenging DoD problems. For example, the Laboratory has helped develop U.S. Army environments for training artificial intelligence agents to defend against unmanned aerial systems, rockets, artillery, and mortar threats. Likewise, the Defense Advanced Research Projects Agency has developed the AlphaDogfight environment to train fully autonomous pilots in aerial combat.

While the development of simulation environments provides an exciting opportunity for training machine learning approaches to DoD mission areas, there are several limitations to current state-of-the-art methods that will likely restrict their integration in the near future. These limitations fall into three related categories: sample efficiency, adaptability, and interpretability. For most environments, DRL methods may require years of virtual simulation time to master a specific task. Additionally, while DRL methods may find a solution to a challenging task, the learned solution is often brittle and not adaptable to changing objectives and perturbations. Finally, the learned solutions are often black box models that are difficult to interpret in any meaningful way, making it a challenge to verify their safety and performance and incorporate them into any composite system or teaming environment. One application which stresses these limitations in particular is human-machine teaming, where agents do not have sufficient samples to train against specific human tendencies, must adapt to changing objectives, and must act in a reasonable or interpretable way to effectively coordinate with their human counterparts.

The MeRLin program is looking to solve complex planning and coordination problems across a range of Department of Defense mission areas. Shown above is a simulation environment known as army GroundForce, developed for training artificial intelligence systems for short-range air defense missions.

The Line-funded program Mission-ready reinforcement learning (MeRLin) is focusing on developing and training DRL algorithms capable of maintaining performance on complex tasks with human allies. To achieve this goal and to help overcome the limitations described above, the program team is researching cross-cutting algorithm solutions. The program staff's initial focus has been to examine fully autonomous teams in complex coordination tasks, where each agent of a team independently controls itself. AI technologies and support their deployment to DoD missions.

MIT LINCOLN LABORATORY SUPERCOMPUTING CENTER