Army advances learning capabilities of drone swarms

August 10 2020

Credit: The Army Research Laboratory

Army researchers developed a reinforcement learning approach that will allow swarms of unmanned aerial and ground vehicles to optimally accomplish various missions while minimizing performance uncertainty.

Swarming is a method of operations where multiple autonomous systems act as a cohesive unit by actively coordinating their actions.

Army researchers said future multi-domain battles will require swarms of dynamically coupled, coordinated heterogeneous mobile platforms to overmatch enemy capabilities and threats targeting U.S. forces.

The Army is looking to swarming technology to be able to execute time-consuming or dangerous tasks, said Dr. Jemin George of the U.S. Army Combat Capabilities Development Command's Army Research Laboratory.

"Finding optimal guidance policies for these swarming vehicles in real-time is a key requirement for enhancing warfighters' tactical situational awareness, allowing the U.S. Army to dominate in a contested environment," George said.

Reinforcement learning provides a way to optimally control uncertain agents to achieve multi-objective goals when the precise model for the agent is unavailable; however, the existing reinforcement learning schemes can only be applied in a centralized manner, which requires pooling the state information of the entire swarm at a central learner. This drastically increases the computational complexity and communication requirements, resulting in unreasonable learning time, George said.

In order to solve this issue, in collaboration with Prof. Aranya Chakrabortty from North Carolina State University and Prof. He Bai from Oklahoma State University, George created a research effort to tackle the large-scale, multi-agent reinforcement learning problem. The Army funded this effort through the Director's Research Award for External Collaborative Initiative, a laboratory program to stimulate and support new and innovative research in collaboration with external

partners.

The main goal of this effort is to develop a theoretical foundation for data-driven optimal control for large-scale swarm networks, where control actions will be taken based on low-dimensional measurement data instead of dynamic models.

The current approach is called Hierarchical Reinforcement Learning, or HRL, and it decomposes the global control objective into multiple hierarchies—namely, multiple small group-level microscopic control, and a broad swarm-level macroscopic control.

"Each hierarchy has its own learning loop with respective local and global reward functions," George said. "We were able to significantly reduce the learning time by running these learning loops in parallel."
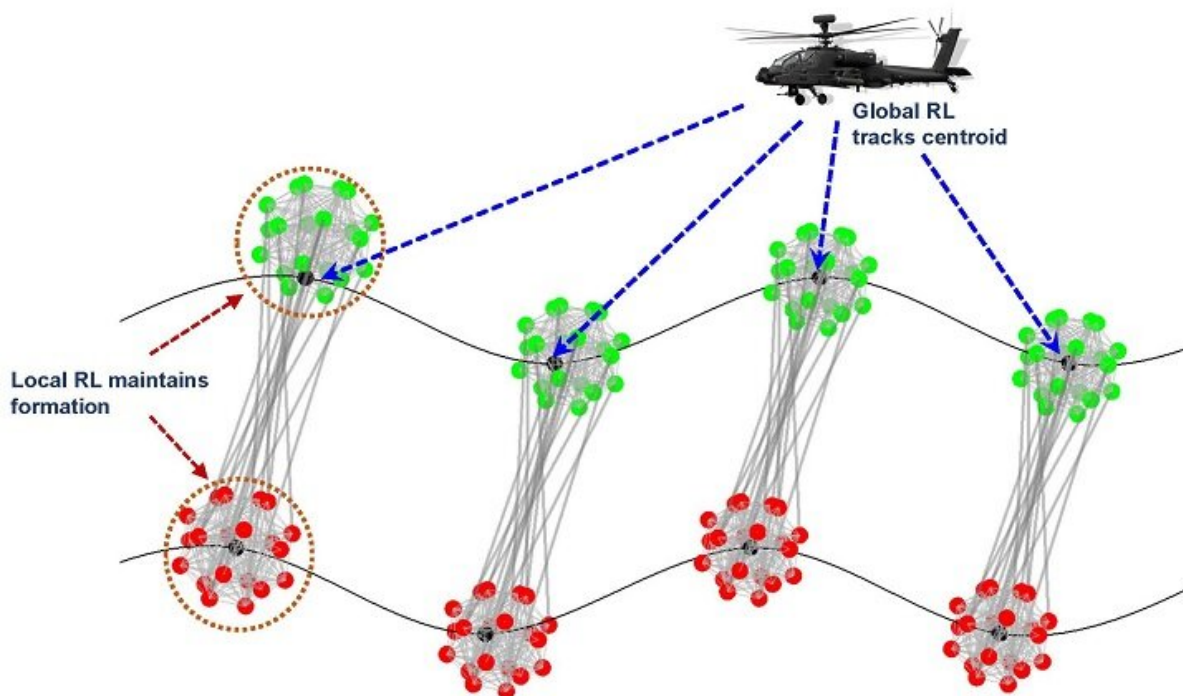
Army researchers envision a hierarchical control for ground vehicle and air vehicle coordination. Credit: U.S. Army graphic

According to George, online reinforcement learning control of swarm boils down to solving a large-scale algebraic matrix Riccati equation using system, or swarm, input-output data.

The researchers' initial approach for solving this large-scale matrix Riccati equation was to divide the swarm into multiple smaller groups and implement group-level local reinforcement learning in parallel while executing a global reinforcement learning on a smaller dimensional compressed state from each group.

Their current HRL scheme uses a decupling mechanism that allows the team to hierarchically approximate a solution to the large-scale matrix equation by first solving the local reinforcement learning problem and then synthesizing the global control from local controllers (by solving a least squares problem) instead of running a global reinforcement learning on the aggregated state. This further reduces the learning time.

Experiments have shown that compared to a centralized approach, HRL was able to reduce the learning time by 80% while limiting the optimality loss to 5%.

"Our current HRL efforts will allow us to develop control policies for swarms of unmanned aerial and ground vehicles so that they can optimally accomplish different mission sets even though the individual dynamics for the swarming agents are unknown," George said.

George stated that he is confident that this research will be impactful on the future battlefield, and has been made possible by the innovative

collaboration that has taken place.

"The core purpose of the ARL science and technology community is to create and exploit scientific knowledge for transformational overmatch," George said. "By engaging external research through ECI and other cooperative mechanisms, we hope to conduct disruptive foundational research that will lead to Army modernization while serving as Army's primary collaborative link to the world-wide scientific community."

The team is currently working to further improve their HRL control scheme by considering optimal grouping of agents in the swarm to minimize computation and communication complexity while limiting the optimality gap.

They are also investigating the use of deep recurrent neural networks to learn and predict the best grouping patterns and the application of developed techniques for optimal coordination of autonomous air and ground vehicles in Multi-Domain Operations in dense urban terrain.

George, along with the ECI partners, recently organized and chaired an invited virtual session on Multi-Agent Reinforcement Learning at the 2020 American Control Conference, where they presented their research findings.

Provided by The Army Research Laboratory

Citation: Army advances learning capabilities of drone swarms (2020, August 10) retrieved 10

April 2024 from