Reinforcement Learning and Numerical Optimization for Air-Ground Robot Team Mission Planning Using High Performance Computing

About DEVCOM ARL:

DEVCOM ARL, as an integral part of the Army Futures Command, is the Army's foundational research laboratory focused on operationalizing science to ensure overmatch in any future conflict. DEVCOM ARL shapes future concepts with scientific research and knowledge and delivers technology for modernization solutions to win in the future operating environment.

RESEARCH LOCATION: Minneapolis, MN

PROJECT DESCRIPTION:

The Army is interested in deploying teams of air and ground robots (i.e. drones and rovers). The air robots are fast and agile but have a short battery life. The ground robots are slower and more restricted but have a much longer battery life. By using the ground robots as mobile charging stations, the air robots can extend their range, and the team becomes a highly effective hybrid system. Mission planning for such systems is quite difficult.

This project will use reinforcement learning to create approximate solutions and use advanced continuous numerical solvers to form plans much faster than mixed-integer linear programming. Both of these approaches require significant computational power, and the plan is to take advantage of the HPCs. The final thrust of the project will be procedural generation and testing of millions of planning.

Under the guidance of a mentor, the intern will gain knowledge in the following activities:

- Develop simulations
- Develop and train reinforcement learning algorithms
- Deploy simulations and algorithms on high-performance computers
- Implement mixed-integer linear programs and other solvers
- Implement and tune heuristic optimization algorithms
- Document and compare results

Intern's Project Schedule:

Week 1: Orientation to University of Minnesota and DEVCOM-ARL labs. Briefing on project status and problem description. Tutorials on mixed-integer linear programming, reinforcement learning, and continuous optimization for path planning.

Week 2: Setup of Git repository and installation scripts. Test runs of algorithms in high-performance computing environment.

Week 3: Development of procedurally generated test scenarios.

Week 4: Implementation of MILP and continuous solvers on batch scenarios on HPCs.

Week 5: Development of reinforcement learning "gym" environment for path planning. Research on RL methods.

Week 6: Training RL on HPCs.

Week 7: Batch testing RL algorithm on HPCs.

Week 8: Comparison of results and brainstorming more complex problems.

Week 9: Re-run algorithms on more complex optimization problems.

Week 10: Document results and transfer code.

ANTICIPATED START DATE:

May/June 2026 – Exact start dates will be determined at the time of selection and in coordination with the selected candidate.

QUALIFICATIONS:

The ideal candidate will have background in software development, Knowledge in robotics and path planning is most relevant. A background in mathematics, mechanical engineering, aerospace engineering or related fields.

ACADEMIC LEVEL:

Degree received within the last 60 months or currently pursuing:

- Bachelor's
- Master's
- Doctoral

DISCIPLINE NEEDED:

- Computer, Information, and Data Science
- Mathematics & Statistics
- Engineering
- Science & Engineering-related