Modeling and Explaining Sequential Behavior

A Series of Unlikely Events and What Will the Robot Do Next?

Understanding sequential behavior is crucial to many defense-related tasks. Why did a drone make a sudden movement away from its destination? Why did a rover choose a certain path? Does a patrolling soldier's route indicate the presence of danger? Two SEI projects offer novel solutions toward modeling and explaining sequential behavior.

Identifying Unlikely Events

Current methods for identifying unlikely or anomalous events require labeled data about what constitutes an unlikely event and the time of human operators to verify predictions. We are using inverse reinforcement learning, an approach based in machine learning, which learns a statistical model of routine and anomalous actions that are taken from each state.

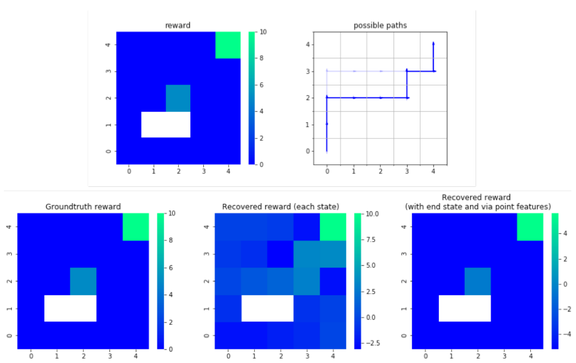

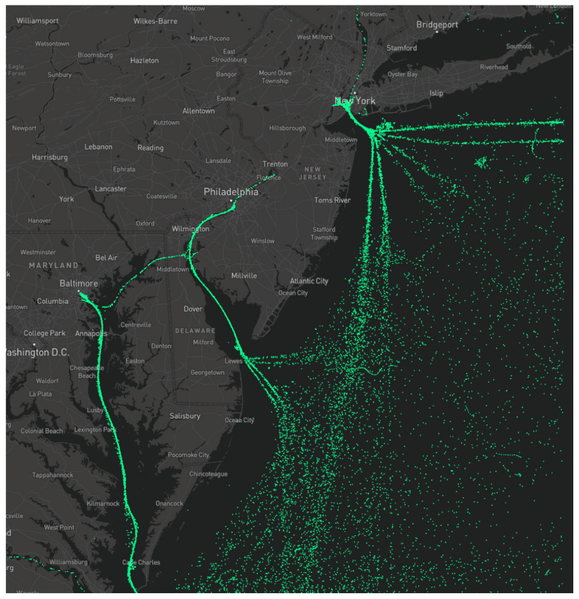

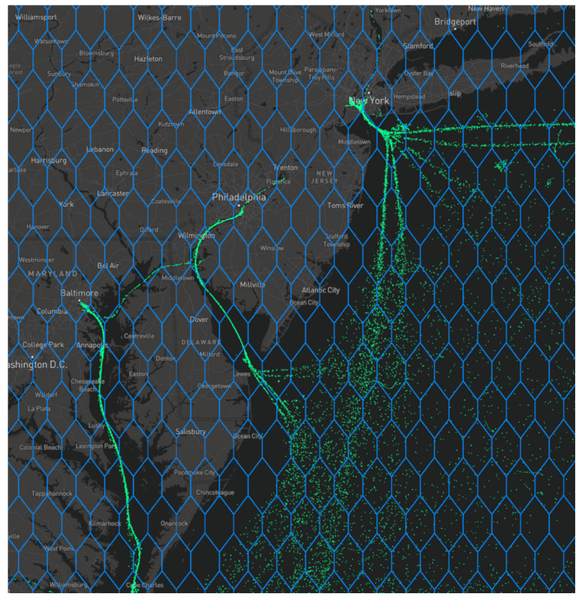

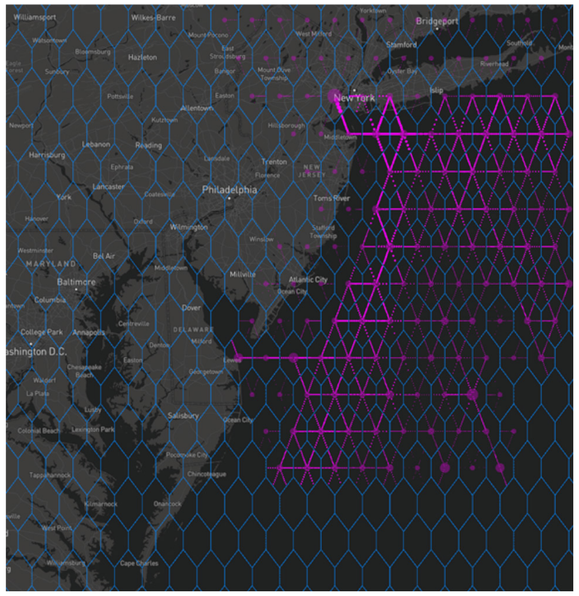

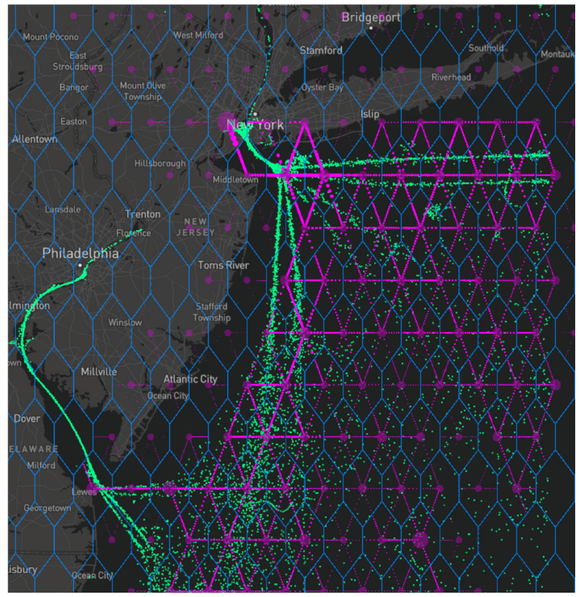

Modeling Ship Paths

Using publicly available Automatic Identification System (AIS) data collected by the U.S. Coast Guard, we use inverse reinforcement learning to model trajectories of marine vessels into New York Harbor. We can use these models to predict where vessels are going, find anomalous behavior, and potentially classify vessel type based on trajectory.

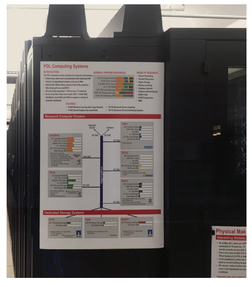

Future Work

In a collaboration with the Carnegie Mellon University Parallel Data Laboratory, we will use inverse reinforcement learning to model behavior of supercomputer users. This collaboration extends our work beyond predicting movements in the physical world and into domains such as cybersecurity, social networks, and more.

Prior Work: Explaining Robot Behavior

For human soldiers working with robot counterparts, being able to predict robot behavior ensures trust and supports human-machine teaming. Our "What Will the Robot Do Next" project has developed algorithms for robots to proactively adapt their behavior to enable users to predict what the robot will do next.

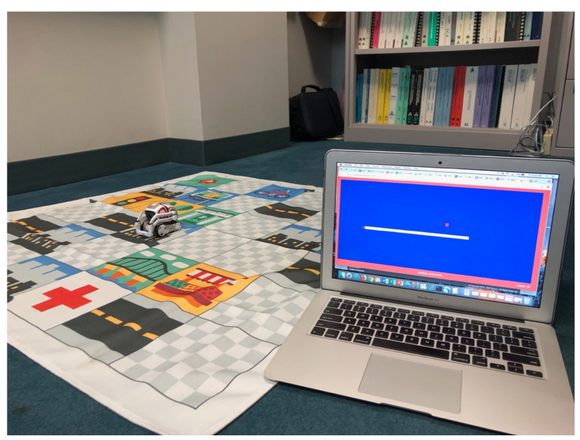

In an ongoing experiment, we are working to predict what people will focus on while performing a dual task: playing a simple video game and observing a Cozmo robot. We will collect dual task data from participants to compare to our predictive models.

Dr. Drew Gifford | sagifford@sei.cmu.edu