

A Series of Unlikely Events

Learning from Sequential Behavior for Activity-Based Intelligence and Modeling Human Expertise

Introduction

Modeling patterns of sequential behavior is a task that underlies numerous difficult artificial intelligence tasks:

- How do I detect when adversaries are deviating from normal routines?
- How can I predict where a ship is going to dock?
- How can I automate the teaching of novice analysts to perform complex tasks as if they were experts?

In this work, we use a class of techniques called **Imitation Learning (IL)** to model sequential behavior to answer questions like these and others.

Methodology

Given observations of behavior:

$$\mathcal{B} = \{((s_1, a_1), (s_2, a_2), \dots)_1, \dots, ((s_1, a_1), \dots)_n\}$$

Learn a policy $\pi: \mathcal{S} \rightarrow \mathcal{A}$ that best explains the behaviors.

Two Kinds of Imitation Learning Algorithms

1. Inverse Reinforcement Learning: Learn a reward function $R: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ that models the preference exhibited in behavior. Then, learn policy π that maximizes expected reward.
2. Behavioral Cloning: Learn π directly to mimic the actions exhibited in the behaviors.

How to use Imitation Learning for...

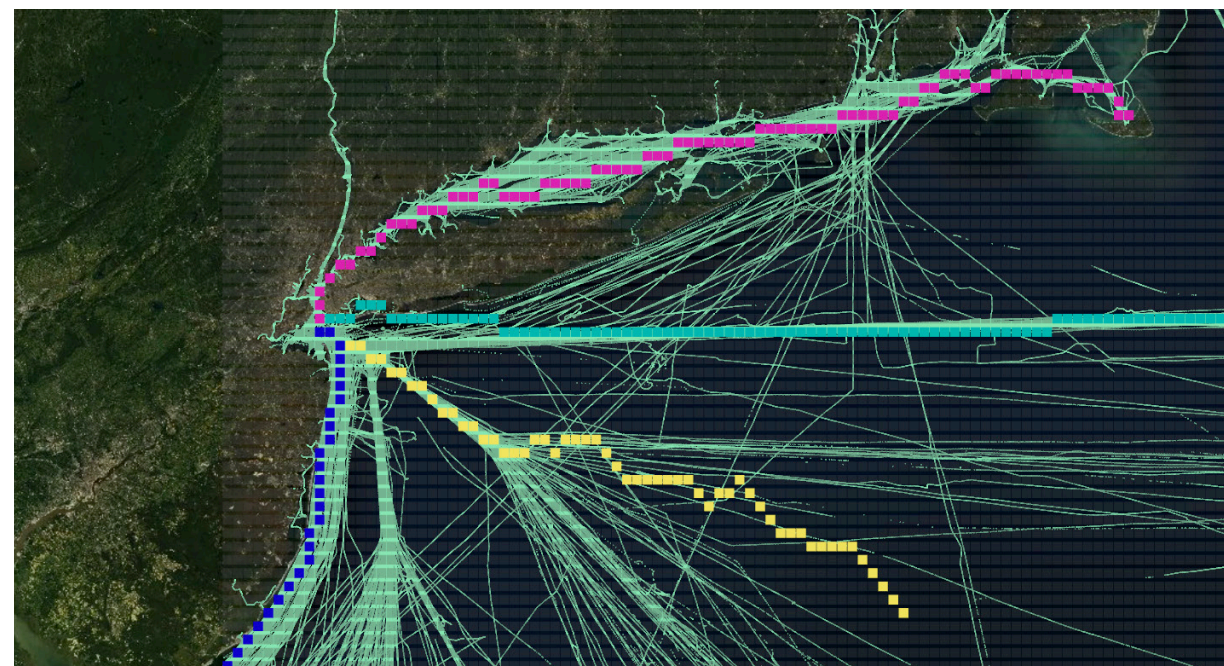
Activity-based Intelligence:

- Learn π from observed behaviors of entities of interest.
- Project future behavior by successively applying π to state.
- Detect anomalous behavior when π deems an action to be of low probability (assumes probabilistic policy).

Teaching expert behavior:

1. Learn π from expert behavior.
2. When a novice is in a state for which she doesn't know the proper action, suggest the one produced by π .

Imitation Learning techniques are an efficient and effective means to perform **activity-based intelligence** or to help automate the education of novices on how to **perform tasks like experts.**



Goals of this work

1. Investigate the practicality (assumptions made, efficiency, scalability, expressiveness) of applying IL to behavioral modeling problems.
2. Apply IL Techniques to DoD/IC relevant problems:
 - Perform efficient implementations that scale to a large number of observations.
 - Build demonstration from data ingestion to visualization tools.
3. Develop techniques that are able to explain, simulate, and demonstrate expert behavior.

Accomplishments

1. Performed technical evaluation of Generative Adversarial Imitation Learning (GAIL) (Ho and Ermon, 2016) and Disturbances for Augmenting Robot Trajectories (DART) (Lee et al., 2017) when applied to modeling nautical vessel behavior.
 - With careful engineering and domain-specific modeling assumptions, we were able to achieve a policy that was able to predict a ship's end destination state within 0.001% of their actual state (technical report forthcoming).
2. Created implementation of Maximum Causal Entropy IRL (MCEIRL) (Ziebart et al., 2010) that is 500x+ faster than academic implementation (to be publicly released).
3. Created demonstration of MCEIRL model applied to U.S. Coast Guard Nautical Vessel Data. (<https://resources.sei.cmu.edu/downloads/IRL-demo>)
4. Developed model with Stephanie Rosenthal (CMU/CSD) and Reid Simmons (CMU/RI) of expert data scientist behavior for the purpose of guiding novice data scientists through challenging tasks (technical report forthcoming).

Ho, Jonathan and Ermon, Stefano. Generative Adversarial Imitation Learning. *Advances in Neural Information Processing Systems (NIPS)*, 29. D. D. Lee et al. (eds). NIPS Foundation. 2016.
 Lee, Jonathan et al. DART: Disturbances for Augmenting Robot Trajectories. *1st Conf. on Robot Learning (CORL) Project*. Nov. 2017.
 Ziebart, Brian D, et al. *Maximum Causal Entropy IRL (MCEIRL)*. School of Computer Science, Carnegie Mellon University. 2010.

Copyright 2020 Carnegie Mellon University.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

DM20-0908