

A Series of Unlikely Events: Learning Patterns by Observing Sequential Behavior

Introduction

Modeling patterns of 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 automate the teaching of novice analysts to perform complex tasks as if they were expert analysts?

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

Methodology

Given observations of behavior:

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

Learn a reward $R: \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ that best explains the observed behaviors.

How IRL algorithms work:

(Two steps)

1. Learn policy $\pi: \mathcal{S} \mapsto \mathcal{A}$ from reward R (policy trained to maximize expected reward as in reinforcement learning).
2. Compute expected behaviors computed from policy π , compared to observed behaviors \mathcal{B} , and use to update reward R .

How to use IRL for...

Activity-based intelligence:

1. Learn R from observed behaviors.
2. If new behavior exhibits low-reward actions in states, flag as abnormal.

Teaching expert behavior:

1. Learn R from expert behavior.
2. When a novice is in a state where she doesn't know the proper action, suggest the one with highest reward.

Inverse Reinforcement Learning techniques are an efficient and effective means to perform **activity-based intelligence** or to teach novices how to **perform tasks like experts.**



©Mapbox, ©OpenStreetMap, and ©DigitalGlobe

Goals of this work

1. Apply IRL Techniques to DoD/IC relevant problems:
 - Employ efficient implementations that scale to a large number of observations.
 - Build demonstration data ingestion to visualization tools.
2. Develop Novel IRL Techniques that are robust to rare events.
3. Develop techniques that are able to explain, simulate, and demonstrate expert behavior.

Progress

1. Implementation of Maximum Causal Entropy IRL that is up to **1000x faster** than academic implementation.

Source: Ziebart, Brian D., J. Andrew Bagnell, and Anind K. Dey. "Modeling interaction via the principle of maximum causal entropy." Proceedings of the 27th International Conference on Machine Learning. Omnipress, 2010.

2. Data ingestion pipeline and visualization of IRL being used to model ship behavior on U.S. Coast Guard Automatic Ship Identification (AIS) data.
3. Investigation into current IRL techniques and how robust they are to rare events, leading to initial formulation of robust IRL technique.
4. Data scientist study to capture expert data scientist behavior.

Additional Figures

RoboschoolAnt-v1 Demonstration

```
[1]: import gym
import roboschool
import matplotlib.pyplot as plt
```

In order to use the roboschool environments, you must import both `gym` and `roboschool`, as many of the methods used to manipulate the environment are in the `gym` library, while much of the functionality specific to roboschool environments are in `roboschool`.

```
[2]: env = gym.make('RoboschoolAnt-v1')
```

The first step is to make the environment by calling `make` with a string identifier that indicates what environment you want to use. Note that when creating an environment in OpenAI gym, you must first "register" it. Roboschool registers an environment when it is imported.

```
[3]: state = env.reset()
#print(state)
```


Copyright 2019 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.

Carnegie Mellon* is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

DM19-1030