A Series of Unlikely Events: Learning from Sequential Behavior for Activity-Based Intelligence and Modeling Human Expertise

Dr. Eric Heim

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.

DM19-1103

An Example of Modeling Behaviors: Ship Movement

Image Credit : ShipsNet Data Set https://github.com/rhammell/shipsnet-detector

What can a model of ship behaviors give you?

What can a model of ship behaviors give you?

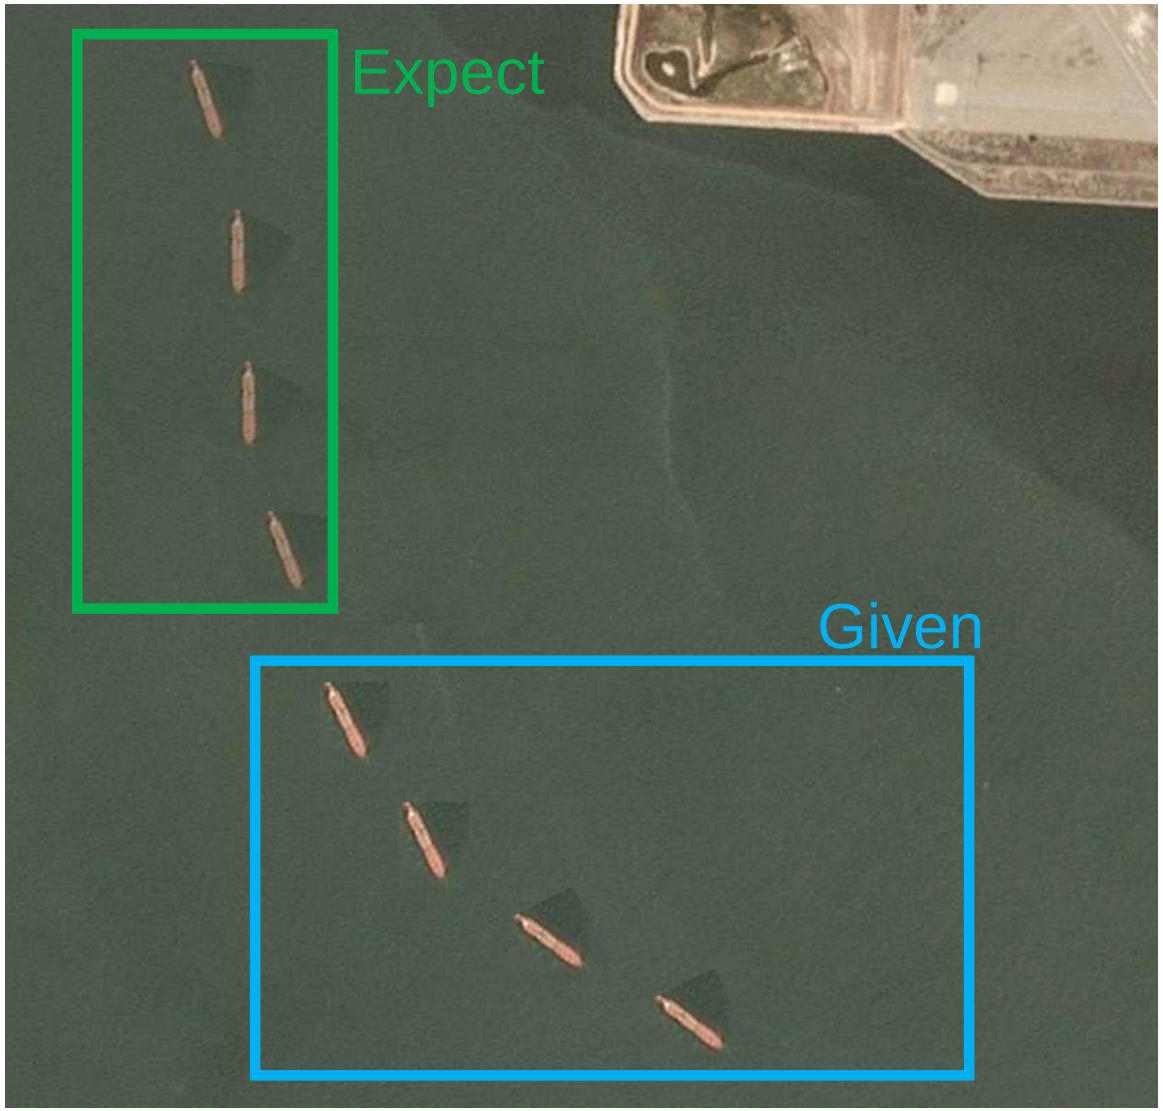

#1

Predictions on where ships are likely to go

What can a model of ship behaviors give you?

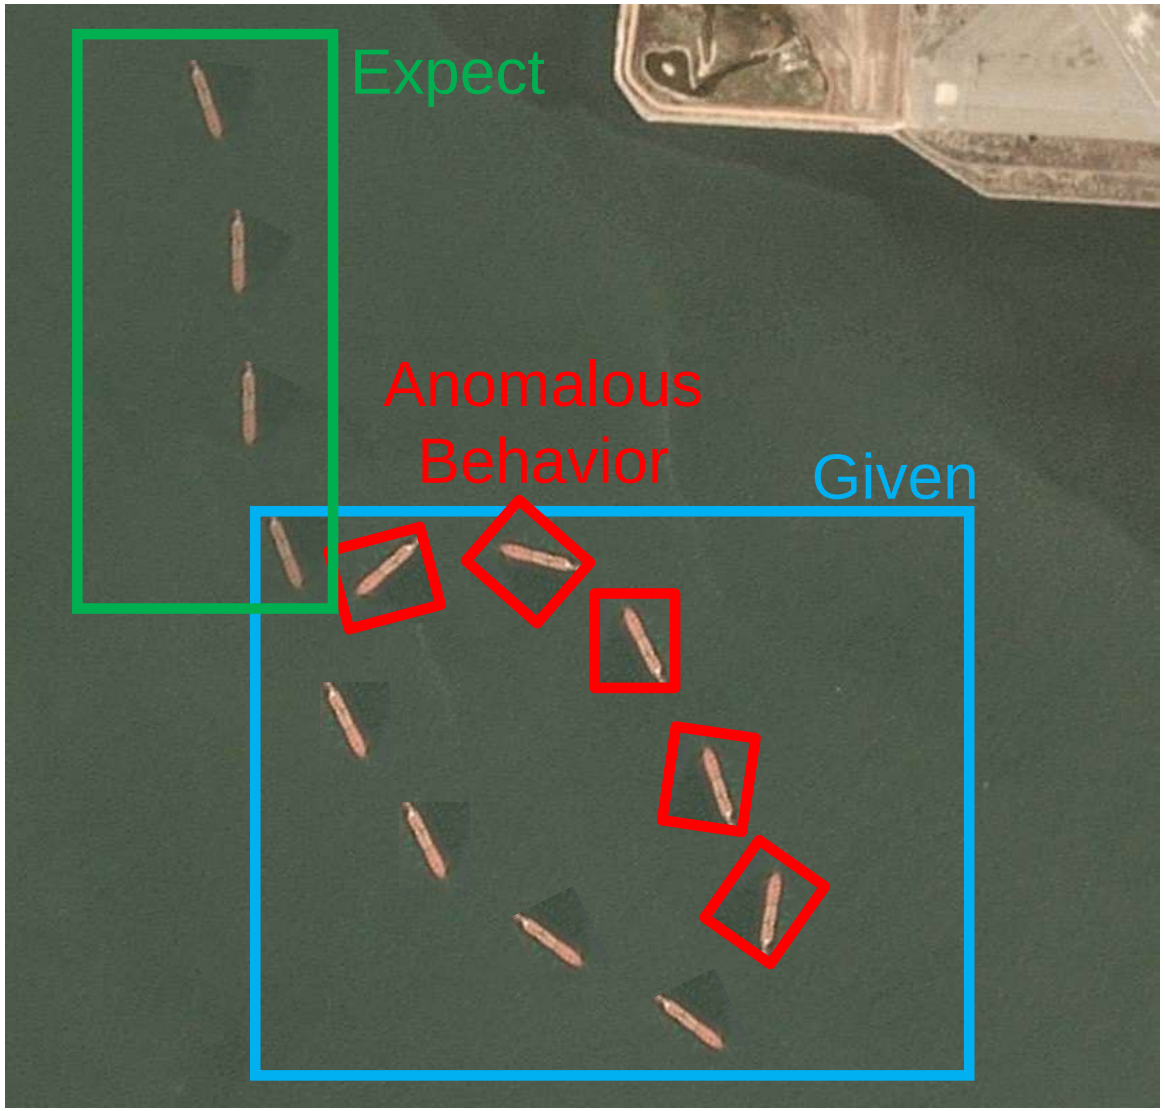

#2

Detections of anomalous ship behavior

What can a model of ship behaviors give you?

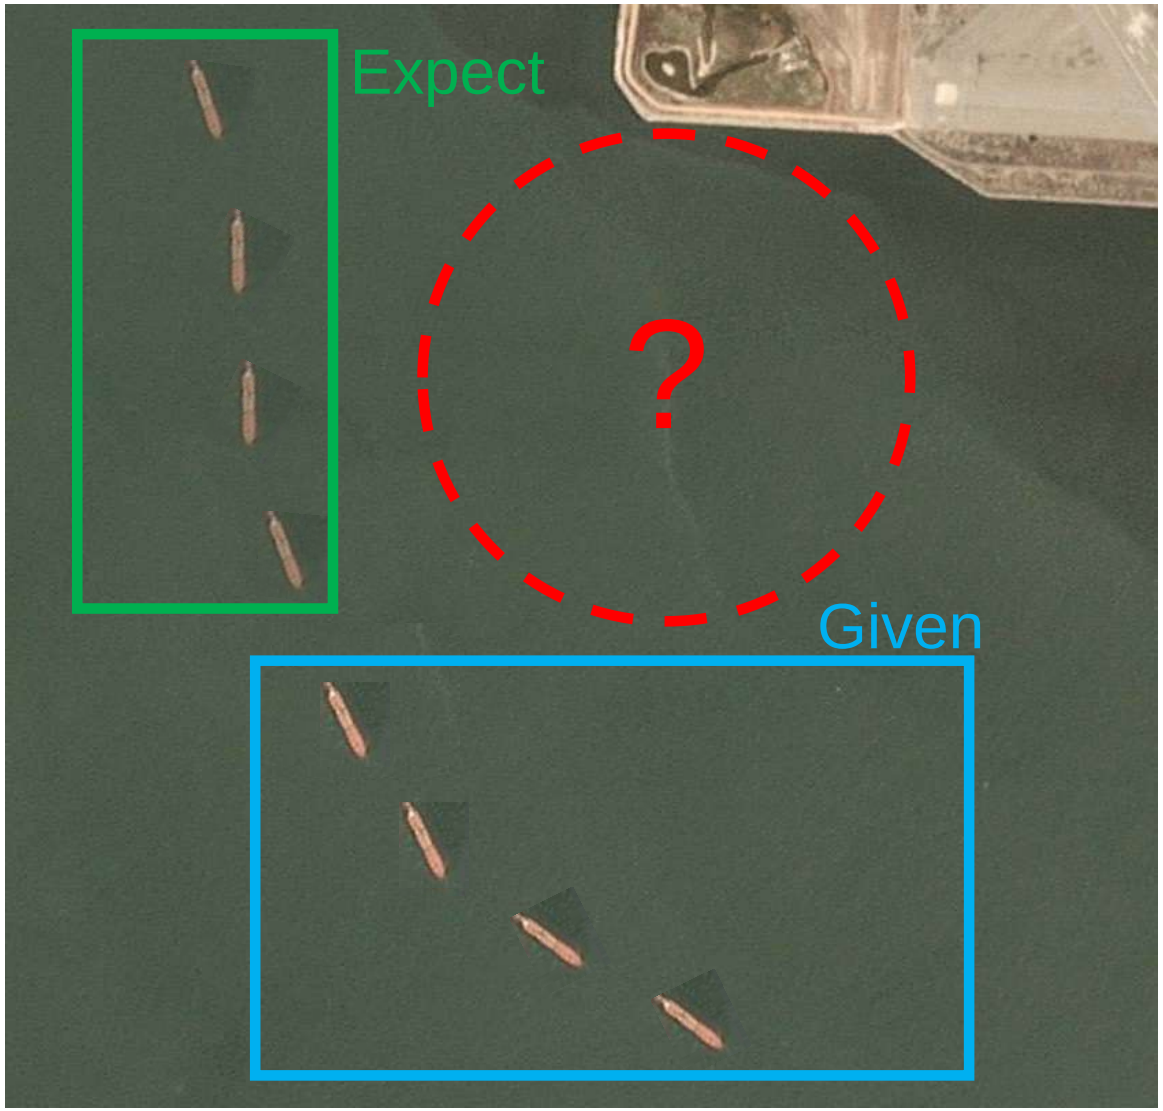

#3

Interferences on areas that ships seem to avoid

What can a model of ship behaviors give you?

#4

Trends of behavior that persist among all ships

In order to model the behaviors typically considered in DoD and IC domains, we require methods that:

1. **Scale** to large data requirements
2. Are **robust** to rare or novel behaviors
3. **Faithfully model the domain** that is considered

The goal of this work is to create **efficient, robust** methods that **faithfully model** behaviors in important operational domains.

Roadmap

- **Background: Inverse Reinforcement Learning (IRL)**
- **Scaling IRL to large problems**
  - **Empirical results**
  - **Demonstration (Coast Guard ship tracking)**
- **Preview of Upcoming Work**
  - **Other Applications of IRL**
  - **Robust IRL**
  - **Modeling human behaviors in IRL**

Roadmap

Background: Inverse Reinforcement Learning (IRL)

Scaling IRL to large problems

Empirical results

Demonstration (Coast Guard ship tracking)

Preview of Upcoming Work

Other Applications of IRL

Robust IRL

Modeling human behaviors in IRL

A Gentle Introduction to Inverse Reinforcement Learning

Main Takeaway: *Inverse Reinforcement Learning (IRL)* takes a set of observed behaviors (captured in data) by one or more agents, and learns the preferences agents have that describe to observed behaviors.

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 function*

$$R: \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$$

That is:

1. High for behavior that is common in the data
2. Low for behavior that isn't

The reward function effectively models preference exhibited in the observed behaviors.

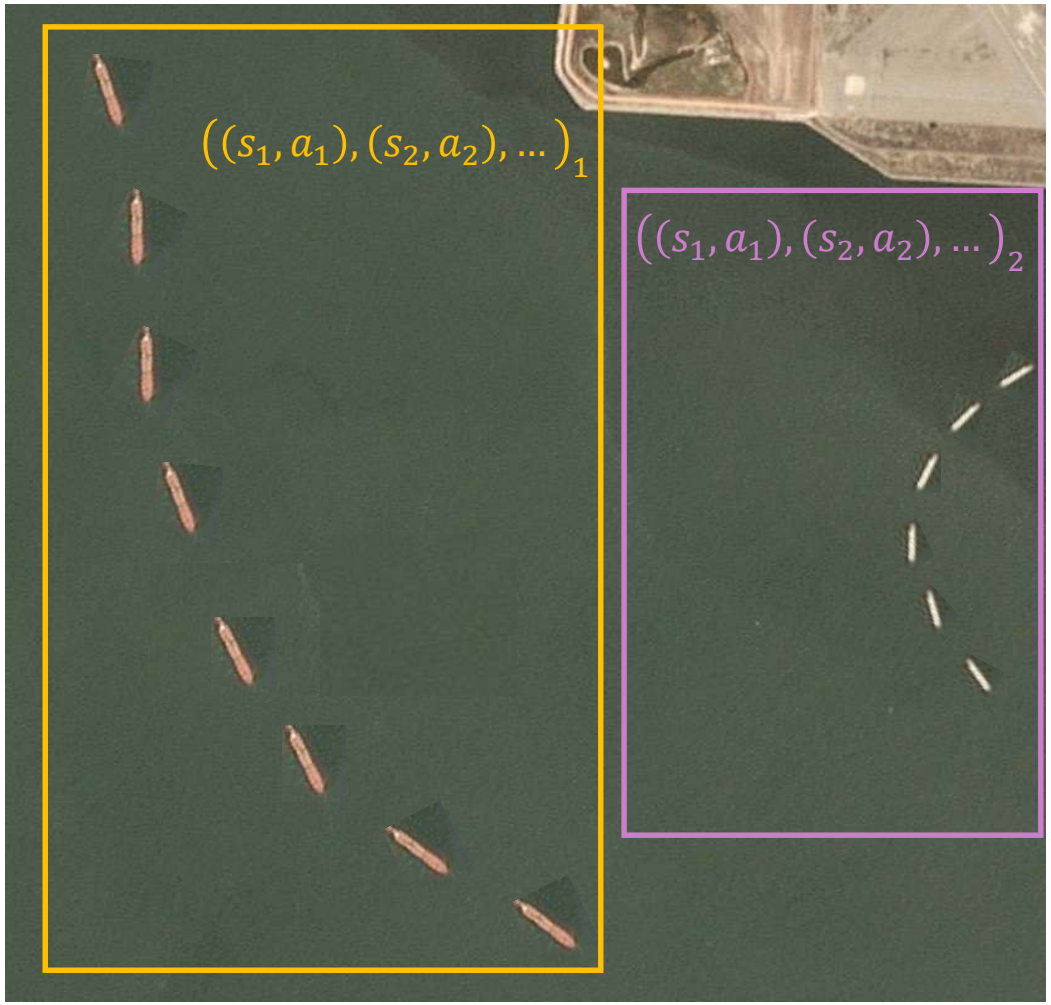

Roadmap

Background: Inverse Reinforcement Learning (IRL)

Scaling IRL to large problems

- **Empirical results**
- **Demonstration (Coast Guard ship tracking)**

Preview of Upcoming Work

- Other Applications of IRL
- Robust IRL
- Modeling human behaviors in IRL

Empirical Results – Computational Efficiency

Requirement: In order for us to apply IRL to large-scale problems, we must write the software that learns a reward function fast.

Environment	# States	# Actions	Sequences	Python	C++	Speedup
10 x 10 Square Grid	100	4	20	3m 15s	5s	40x
25 x 25 Square Grid	625	4	20	58m 24s	1m 10s	50x
100 x 100 Hex Grid	10,000	7	226	23h 51m 34s	3m 7s	459x

Instead of learning a model in a day, our implementation took 3 minutes.


Map Controls
Original Coordinates
Mapped Coordinates
Density Heatmap
Coordinate Tool Tips
mapbox
© Mapbox © OpenStreetMap Improve this map © DigitalGlobe

Roadmap

Background: Inverse Reinforcement Learning (IRL)

Scaling IRL to large problems

- **Empirical results**
- **Demonstration (Coast Guard ship tracking)**

Preview of Upcoming Work

- **Other Applications of IRL**
- **Robust IRL**
- **Modeling human behaviors in IRL**

Other Examples of Behaviors to Model

To identify important events

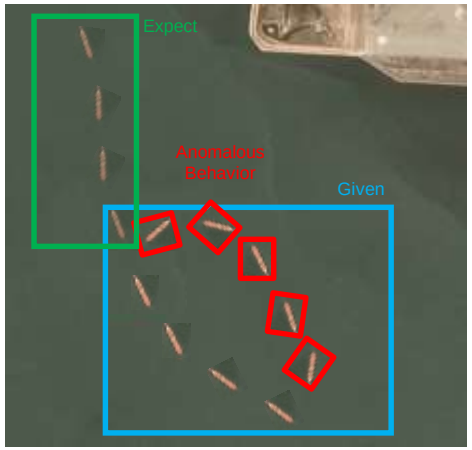

Image Credit : DARPA Mind's Eye Project
http://www.cs.colostate.edu/~draper/MindsEye.php

To teach how to perform tasks

Robust Inverse Reinforcement Learning

Initial Position

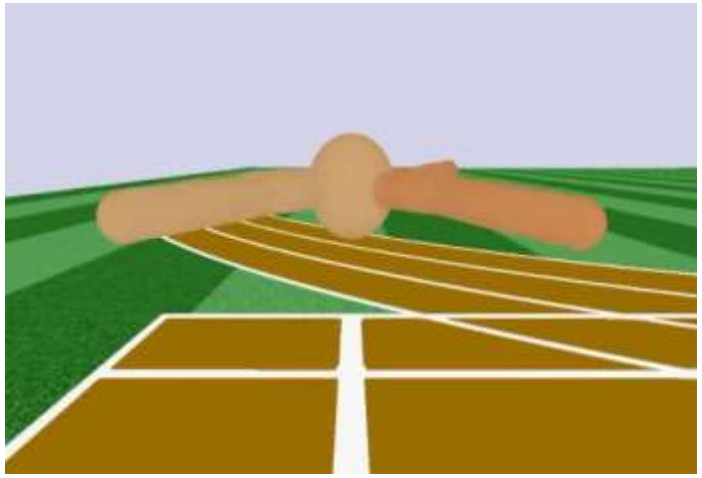

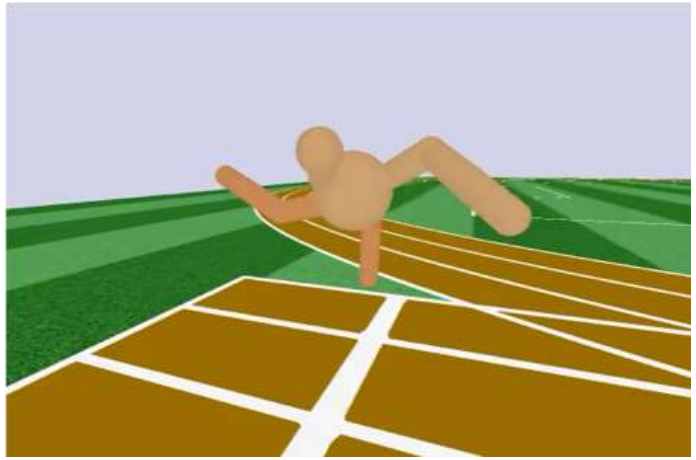

Learned Model

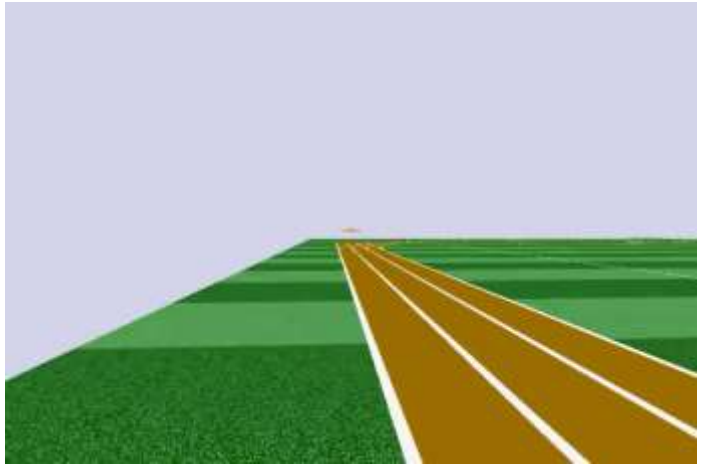

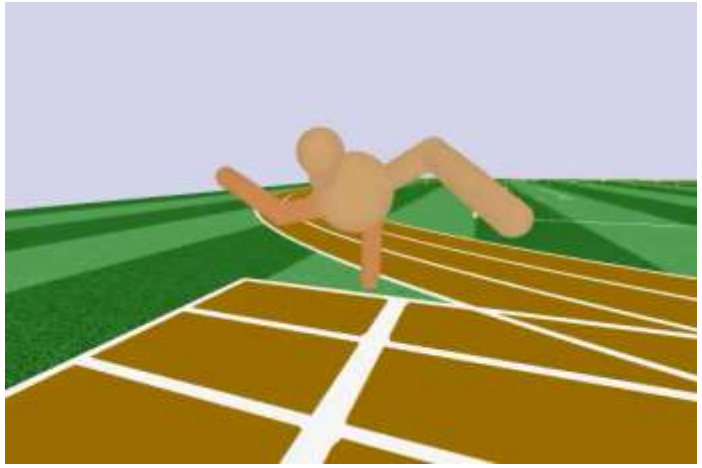

Robust Inverse Reinforcement Learning

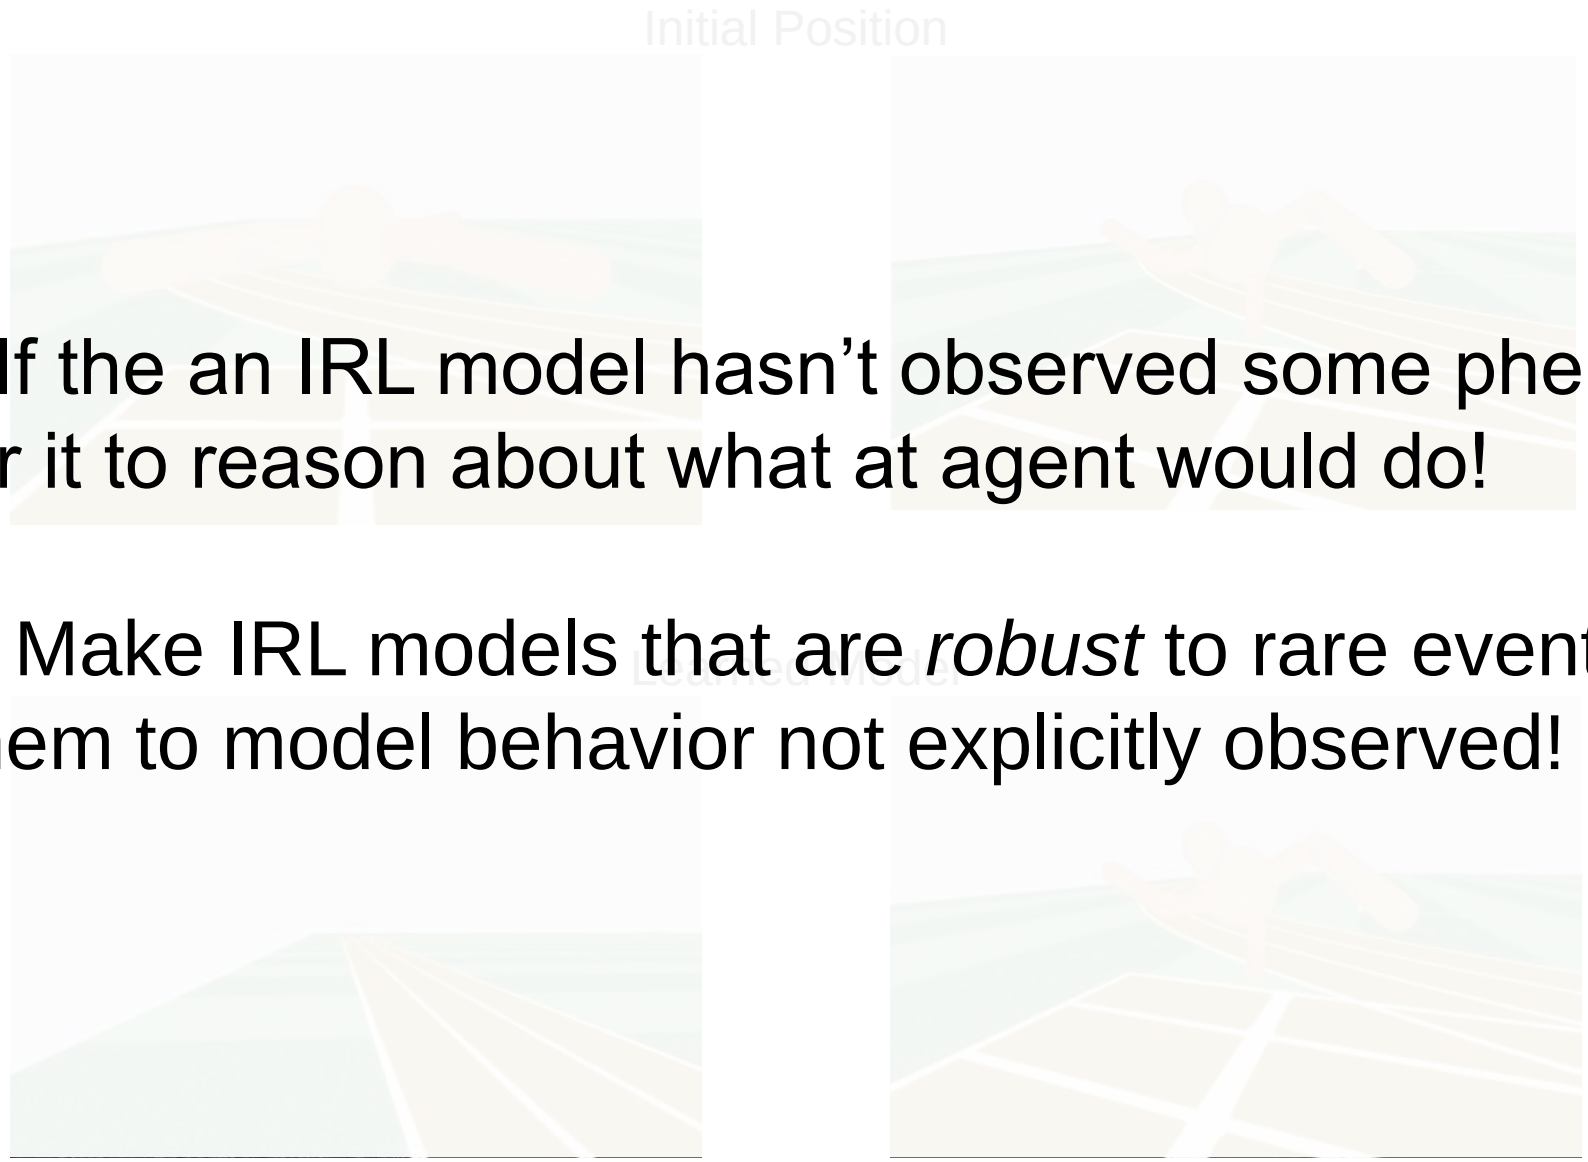

Problem: If the an IRL model hasn't observed some phenomenon, it's hard for it to reason about what at agent would do!

Our Goal: Make IRL models that are *robust* to rare events, allowing them to model behavior not explicitly observed!

Learning Patterns of Human Behavior with IRL

Modeling how day to day activities affect sleep

Dr. Anind Dey
University of Washington

Modeling how expert data scientists explore data

Dr. Stephanie Rosenthal
Carnegie Mellon University

Future Vision

NEAR (~1 year)	MID (2-3 years)	FAR (3-5 years)
Large-scale demonstration of IRL on U.S. Coast Guard shipping data Novel Robust IRL method formalized, developed, rigorously evaluated. All code from the project publicly released. Collaborate with a DoD/IC partner to begin applying IRL to their problem.	Create proof of concept for a sequential behavior model on medium-scale problem in a DoD/IC operational setting: 1. Modeling complex, evolving phenomenon (e.g. network traffic) 2. Modeling expert behavior 3. Others	Deploy a sequential behavior model on large-scale or uniquely challenging problem in a DoD/IC operational setting:

Conclusion

Inverse Reinforcement Learning (IRL) is a class of techniques to model sequential behavior that can be used for:

1. Activity-based intelligence
2. Robotic control
3. Teaching novices how to perform expert tasks

We are:

1. Scaling IRL techniques so they can be applied to large-scale DoD/IC problem domains

   Our work has taken training time from days to minutes.

2. Creating robust IRL techniques that model rare or novel behaviors

   Our work seeks to make models more reliable.

3. Using IRL to model human behaviors for the purpose of teaching complex tasks to novices, and giving tips on how to adjust behavior to live healthier lives.

   We are using models of human behavior to improve quality of life and to teach.

Where you come in:

If you have behaviors that can be modeled from data, we would love to work with you!