A Series of Unlikely Events

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

Eric Heim **etheim@sei.cmu.edu**

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

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.

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 use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

DM20-0830

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

An Example of Modeling Behaviors: Ship Movement

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

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

What can a model of ship behaviors give you?

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

What can a model of ship behaviors give you?

#1

Predictions on where ships are likely to go

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

What can a model of ship behaviors give you?

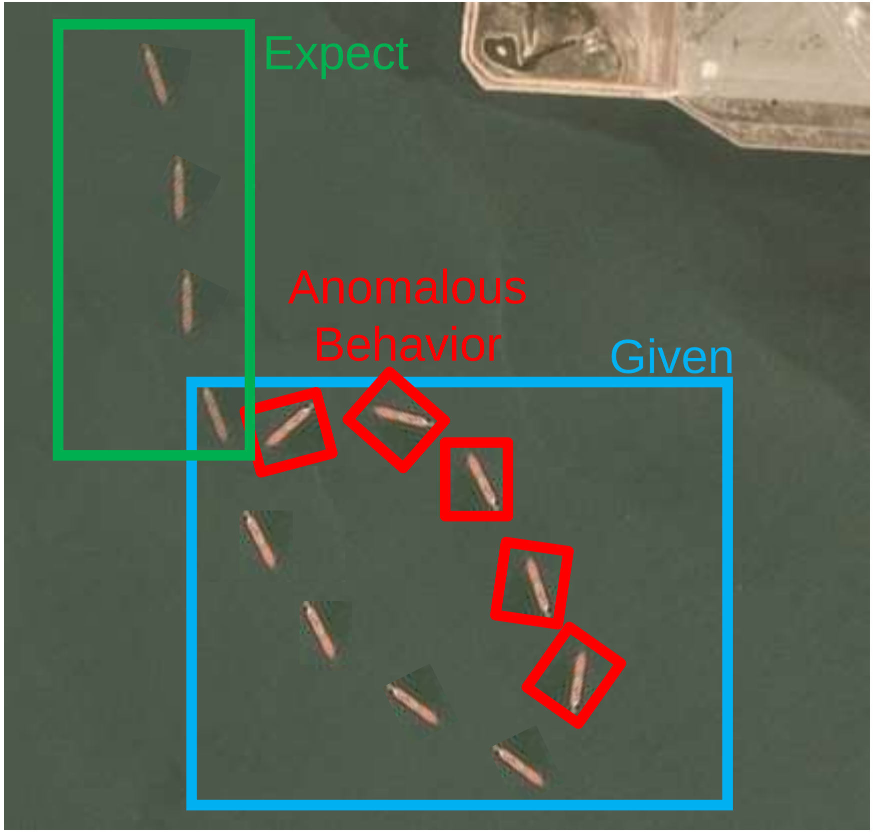

#2

Detections of anomalous ship behavior

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

What can a model of ship behaviors give you?

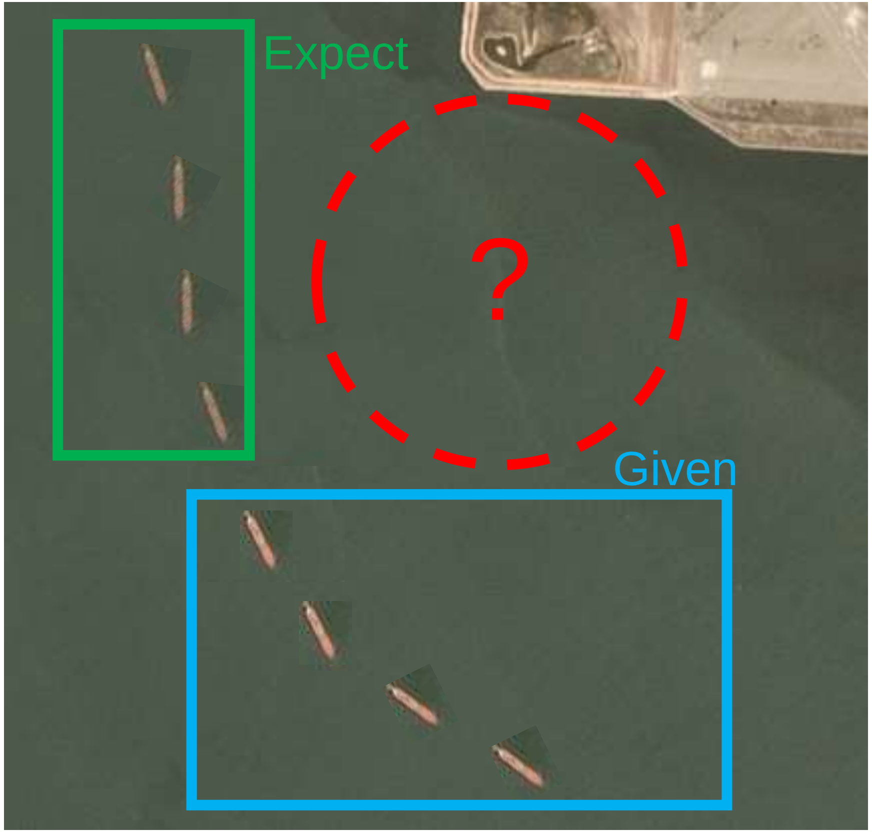

#3

Interferences on areas that ships seem to avoid

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

What can a model of ship behaviors give you?

#4

Trends of behavior that persist among all ships

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Other Examples of Behaviors to Model

To identify important events

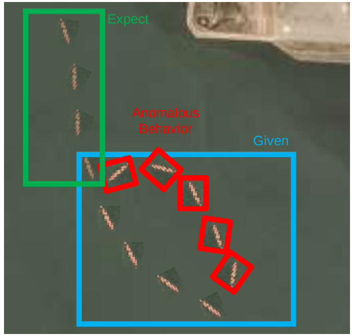

Image Credit : DARPA Mind's Eye Project
http://www.cs.colostate.edu/~draper/MindsEye.php

To teach how to perform tasks

Other Image Credits: www.sei.cmu.edu and www.cmu.edu

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

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

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

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.

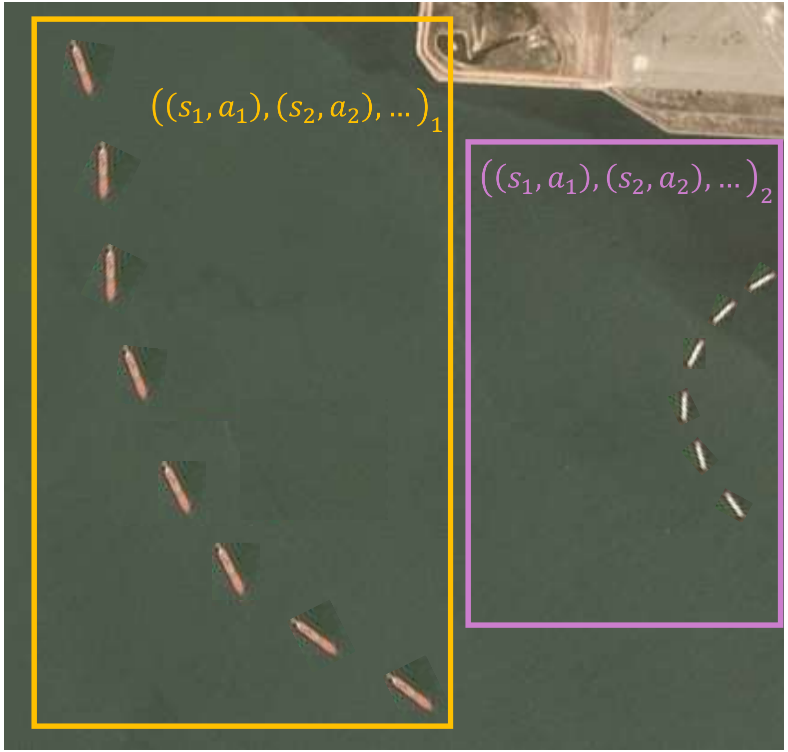

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

IRL in a little more detail

Given: $\mathcal{T} = \left\{ \left((s,a)_1, \dots, (s,a)_{n_1} \right)_1, \dots, \left((s,a)_1, \dots, (s,a)_{n_m} \right)_m \right\}$, where $s \in \mathcal{S}, a \in \mathcal{A}$

Two Steps (Iterate between until convergence):

1. Learn a policy $\hat{\pi}: \mathcal{S} \times \mathcal{A} \mapsto [0,1]$, such that $\forall_{s \in \mathcal{S}} \sum_{a \in D_{\mathcal{A}}(s)} \hat{\pi}(a|s) = 1$
   based on the current estimate of reward $\hat{r}$, using forward reinforcement learning
2. Learn a reward $\hat{r}: \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
   based on how closely expected trajectories from π are to $\mathcal{T}$

Main Ideas:

π is meant to recover some π^* that generated $\mathcal{T}$

r models the underlying reward structure that drives π^*

Both can be applied to states and actions that were not observed!

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Demo

https://resources.sei.cmu.edu/downloads/IRL-demo/

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

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.

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Next Steps: More Rigorous Evaluation

Tasks we want to show IRL can do:

Per ship:

1. Change prediction (change in speed, sharp turn)
2. Destination prediction.
3. 3. Collision detection

Per Region:

1. Affect of season on ship behavior
2. Affect of ship type on ship behavior

Train a MCEIRL (Ziebart; 2010) model on different seasons. Compare policies using KL Divergence.

Base season: Winter 2016
Compared to Spring, Summer, Fall, Winter 2015

$$D_{KL}^{s}(\pi_w, \pi_o) = \sum_{a \in D_{\mathcal{A}}(s)} \pi_w(a|s) \log \frac{\pi_w(a|s)}{\pi_o(a|s)}$$

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Experiment: Divergence between seasonal policies

Winter 2016, Spring 2015

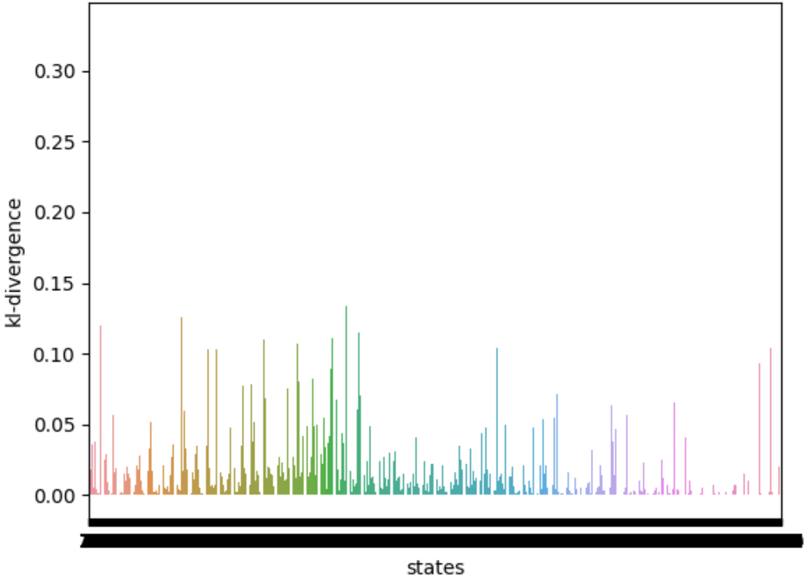

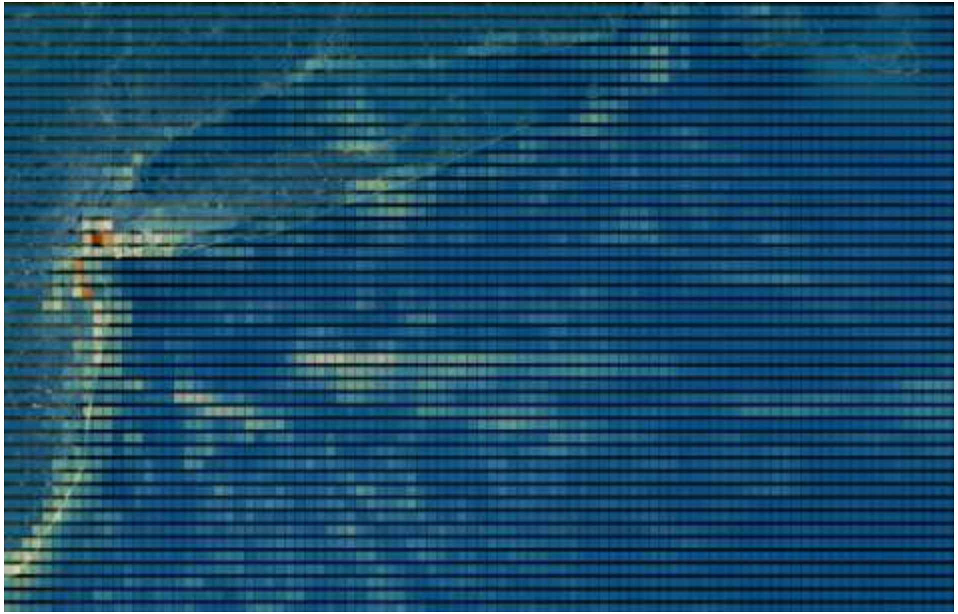

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Experiment: Divergence between seasonal policies

Winter 2016, Fall 2015

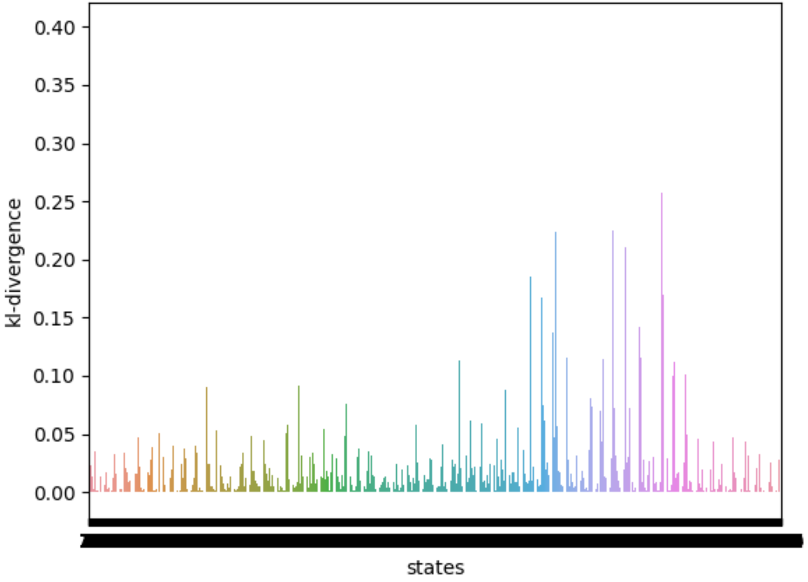

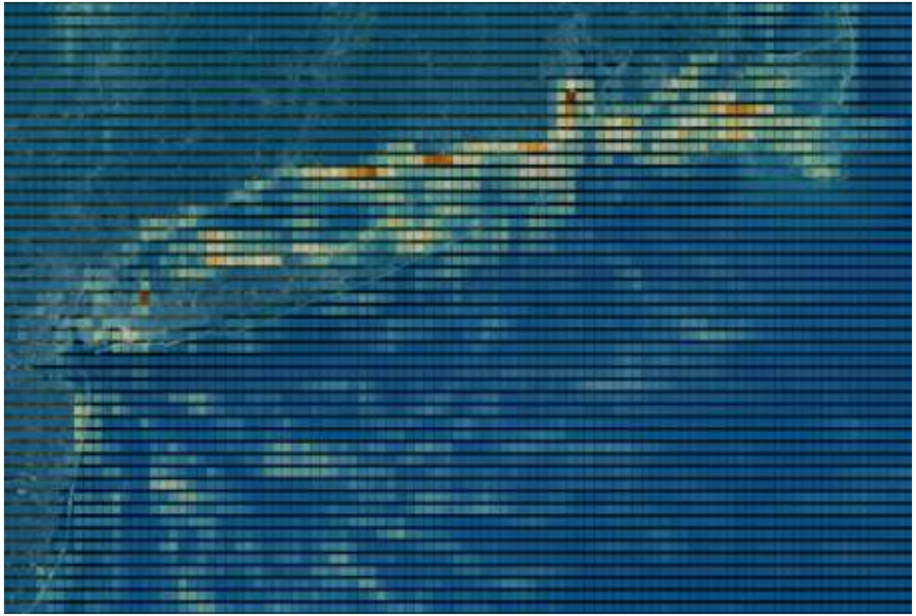

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Experiment: Divergence between seasonal policies

Winter 2016, Summer 2015

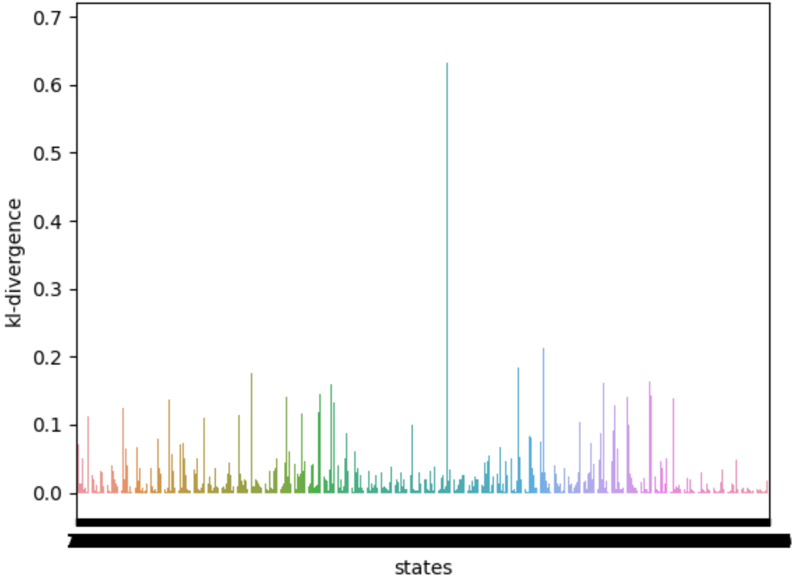

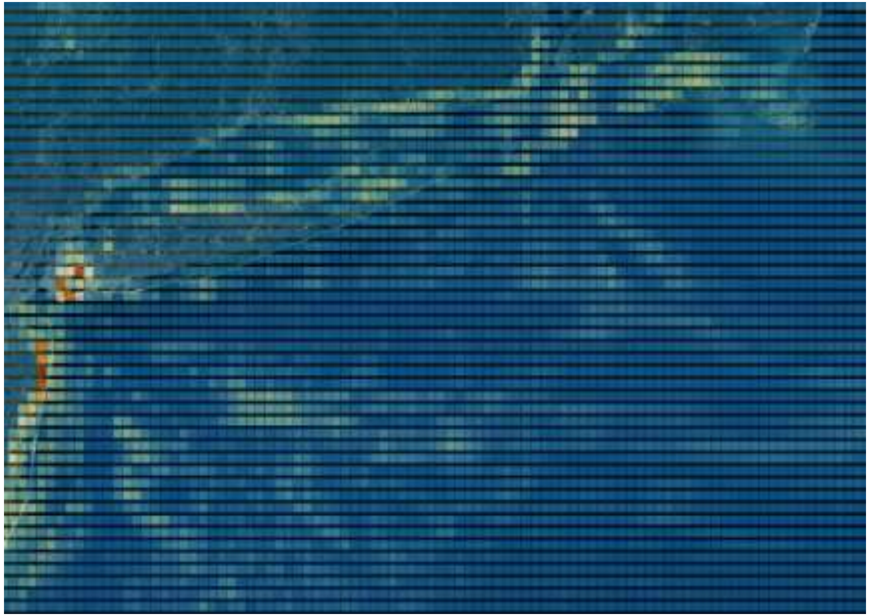

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Experiment: Divergence between seasonal policies

Winter 2016, Summer 2015

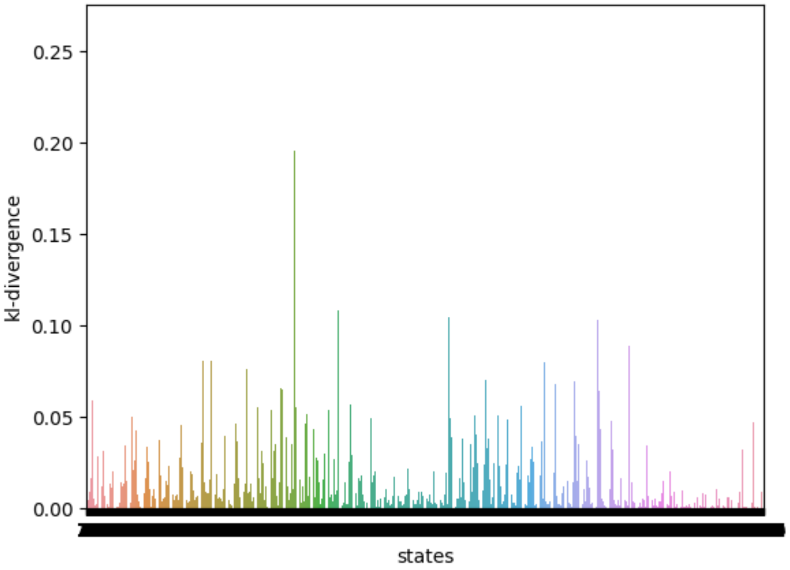

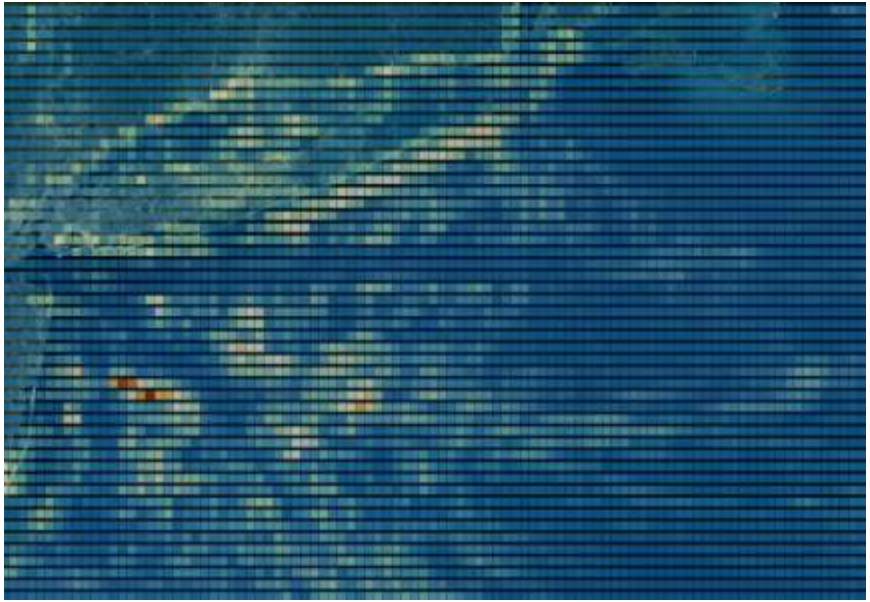

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

More Ongoing Work: More Expressive State/Action Spaces

In the demonstration…

1. States $\mathcal{S}$ are discrete "boxes" in the ocean
2. Actions $\mathcal{A} = \{North, Northeast, East, Southeast, \ldots\}$

Nice for a simple demonstration, but there's more information to represent a ship's state.

U.S. Coast Guard Automatic Identification System Data		
MMSI Identifier	IMO Identifier	Date/Time
LAT	LON	SOG
COG	Heading	Callsign
Vessel Type	Status (activity)	Length
Width	Draft	Cargo (Hazardous)

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Ongoing Work: More Expressive State/Action Spaces

In the demonstration…

1. States $\mathcal{S}$ are discrete "boxes" in the ocean
2. Actions $\mathcal{A} = \{North, Northeast, East, Southeast, \ldots\}$

Nice for a simple demonstration, but there's more information to represent a ship's state.

For the Demo…

U.S. Coast Guard Automatic Identification System Data		
MMSI Identifier	IMO Identifier	Date/Time (discretized)
LAT (discretized)	LON (discretized)	SOG
COG	Heading	Callsign
Vessel Type	Status (activity)	Length
Width	Draft	Cargo (Hazardous)

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Ongoing Work: More Expressive State/Action Spaces

In the demonstration…

1. States $\mathcal{S}$ are discrete "boxes" in the ocean
2. Actions $\mathcal{A} = \{North, Northeast, East, Southeast, \ldots\}$

Nice for a simple demonstration, but there's more information to represent a ship's state.

Now…

U.S. Coast Guard Automatic Identification System Data		
MMSI Identifier	IMO Identifier	Date/Time (discretized)
LAT	LON	SOG
COG	Heading	Callsign
Vessel Type	Status (activity)	Length
Width	Draft	Cargo (Hazardous)

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Ongoing Work: More Expressive State/Action Spaces

In the demonstration…

1. States $\mathcal{S}$ are discrete "boxes" in the ocean
2. Actions $\mathcal{A} = \{North, Northeast, East, Southeast, \ldots\}$

Nice for a simple demonstration, but there's more information to represent a ship's state.

Now…

U.S. Coast Guard Automatic Identification System Data		
MMSI Identifier	IMO Identifier	Date/Time (discretized)
LAT	LON	SOG
COG	Heading	Callsign
Vessel Type	Status (activity)	Length
Width	Draft	Cargo (Hazardous)

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Preliminary Results: Some Learned Ship Behaviors

Generative Adversarial Imitation Learning (GAIL) (Ho and Ermon; 2016)

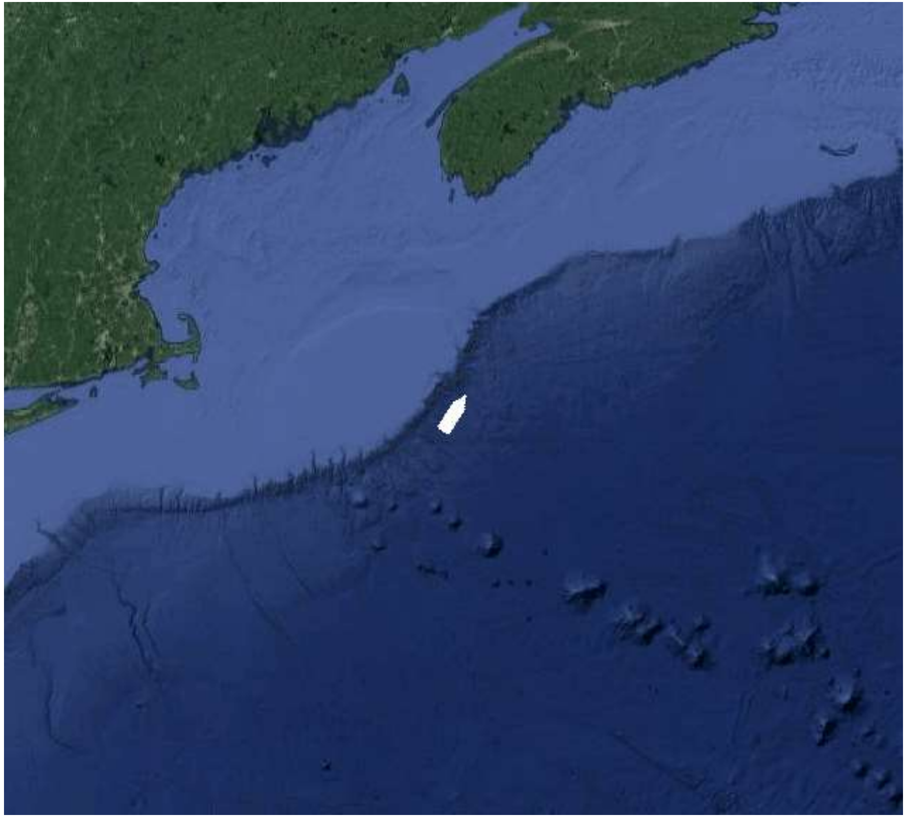

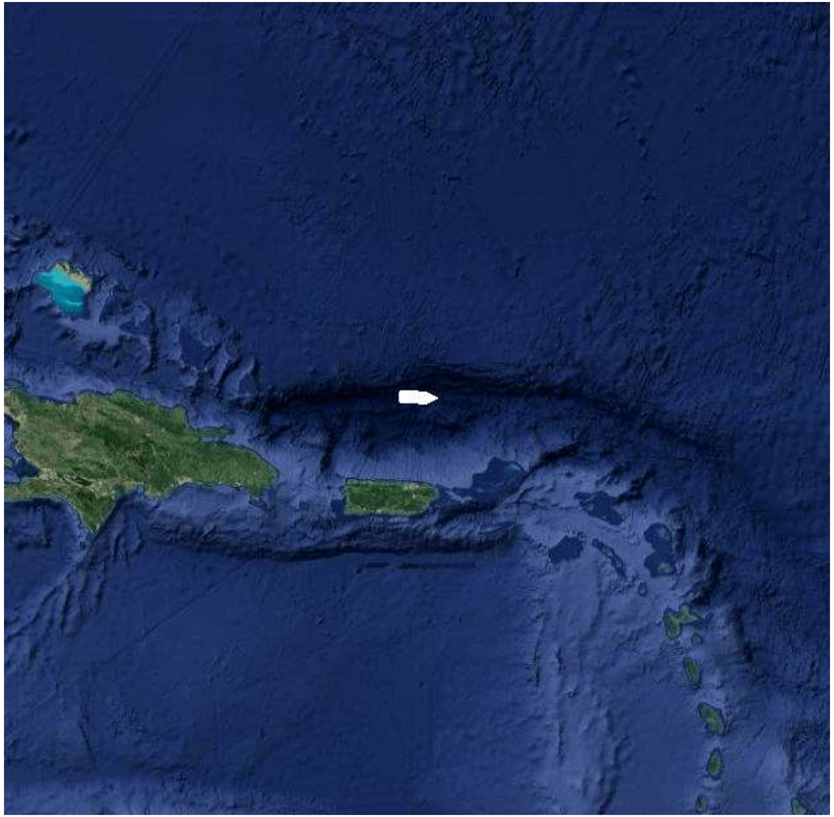

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Preliminary Results: Some Learned Ship Behaviors

Generative Adversarial Imitation Learning (GAIL) (Ho and Ermon; 2016)

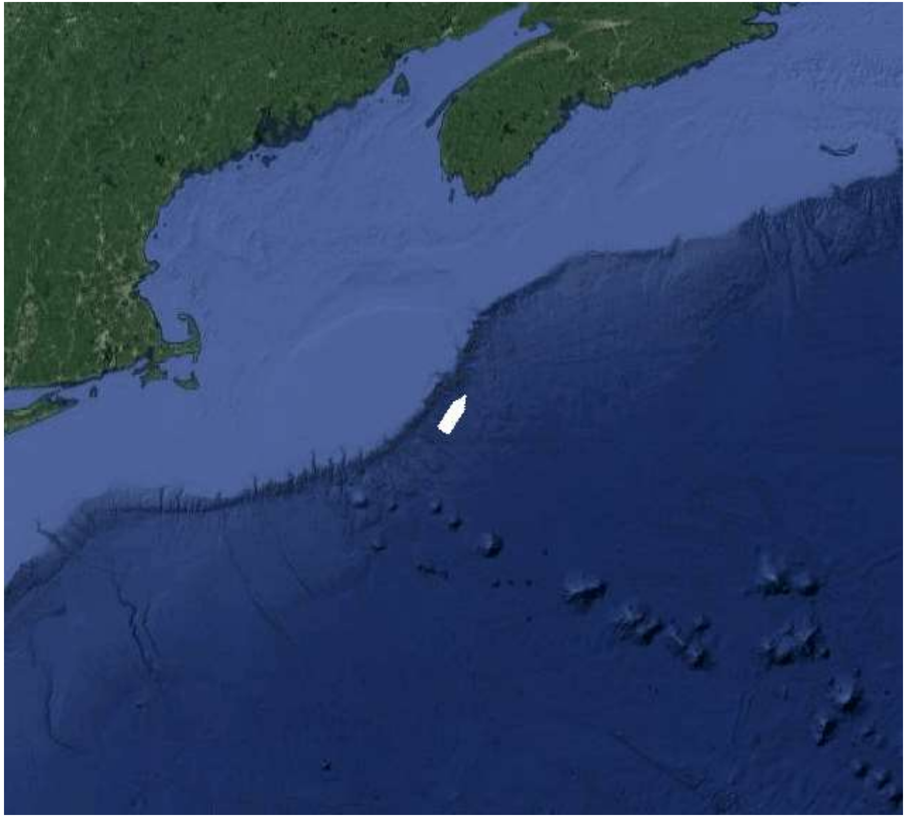

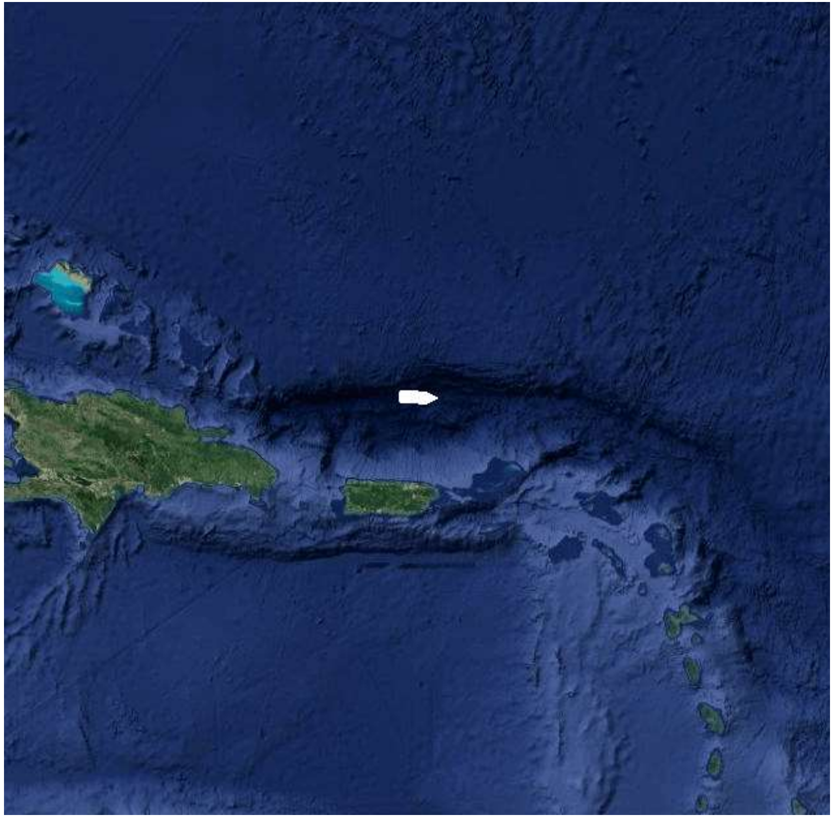

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Modeling Expert Behavior

Main Idea: Capture experts doing sequential tasks; Learn a policy that can predict expert behavior

Potential Applications in Intelligence Analysis:
1. Create tools to teach novice analysts
2. Integrate policy in software toolchain for recommendations

Today: Modeling expert data scientist behavior

Stephanie Rosenthal

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Modeling Expert Data Scientist Behavior

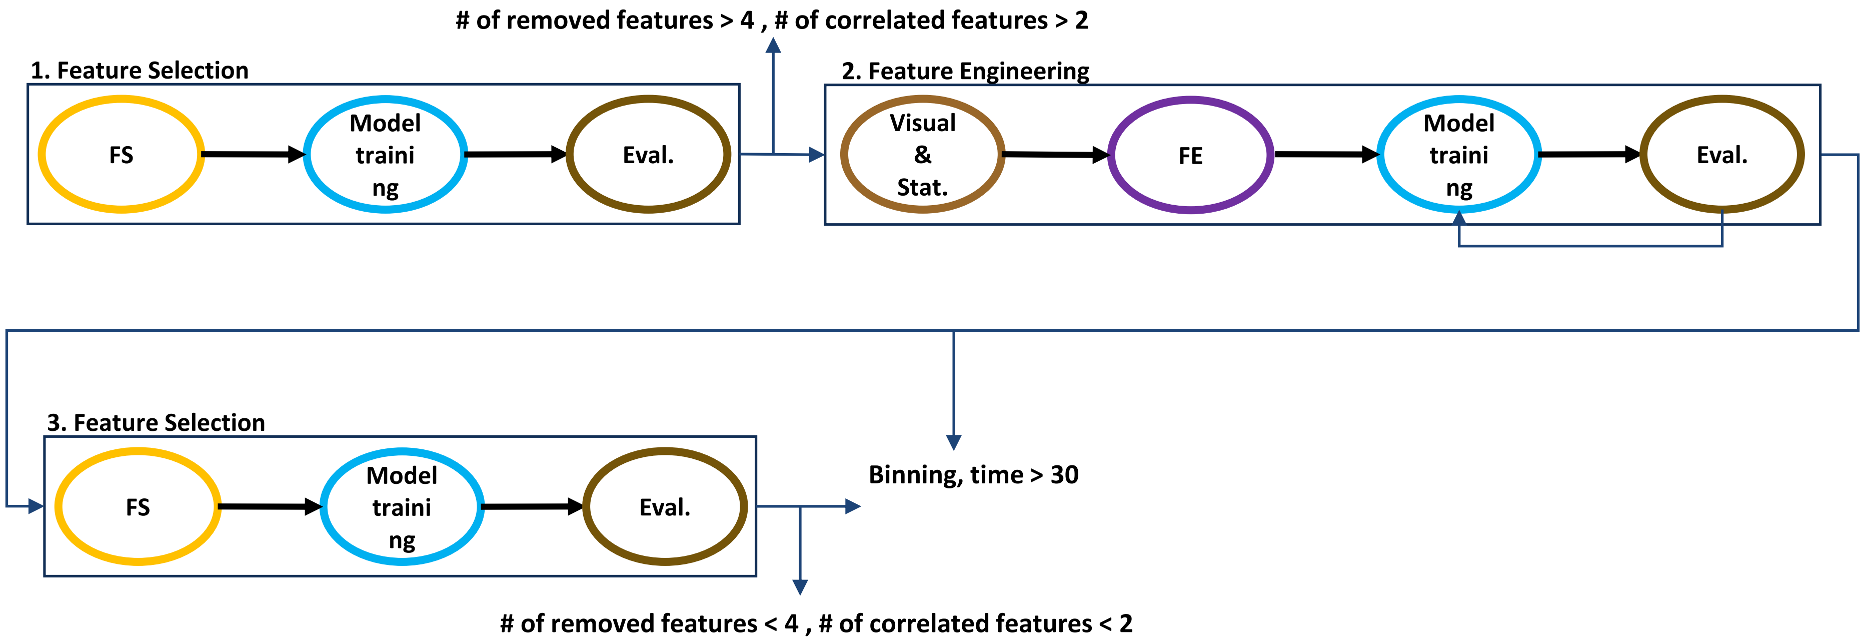

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Code Translation

```
df3 = df2[['WindGustSpeed','Rainfall','Pressure3pm','target']]
```

```
X_train, X_test, y_train, y_test = data_split_train_and_test(df3,0.3)
```

```
columns_to_scale = ['WindGustSpeed','Rainfall','Pressure3pm']
X_train, X_test = feature_scaling(X_train, X_test,columns_to_scale,'MinMaxScaler')
```

```
classifier = train_model('RandomForest',X_train,y_train, {'n_estimators':300, 'min_samples_leaf':20})
eval_get_cm(classifier, X_test, y_test)
```

`df3 = df2[['WindGustSpeed','Rainfall','Pressure3pm','target']]`	`feature_selection`
`df2.dtypes`	`describe_dataframe`
`df3 = FE_divide_numeric_feature_to_ranges(df3,'WindGustSpeed',10)`	`FE_binning`
`X_train, X_test, y_train, y_test = data_split_train_and_test(df3,0.3)`	`splitting`
`classifier = train_model('RandomForest',X_train,y_train, {'n_estimators':300, 'min_samples_leaf':20})`	`train`
`eval_get_cm(classifier, X_test, y_test)`	`eval`

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

From Code to Probabilities of Taking an Action

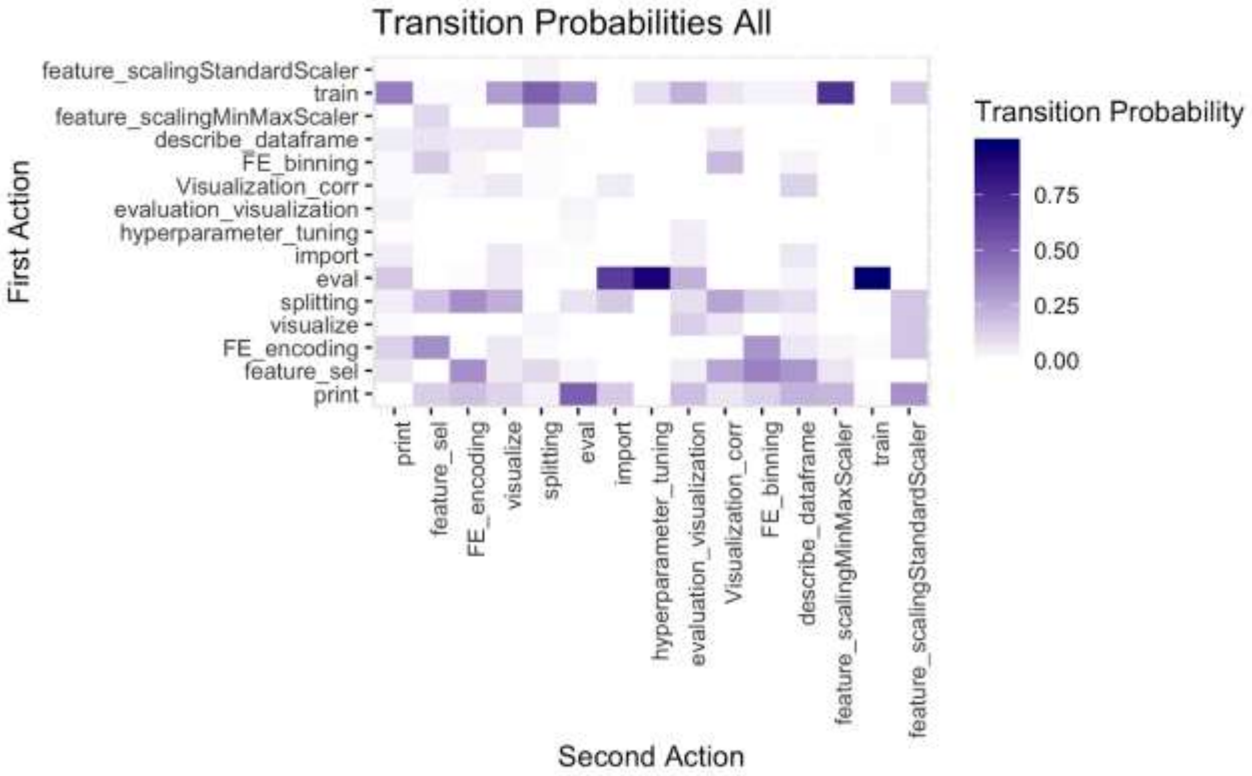

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

From Code to Probabilities of Taking an Action

Log Likelihood for each task

	Adult					Telecom					Rain					
	P1	P2	P3	P4	SUM	P1	P2	P3	P4	SUM	P1	P2	P3	P4	SUM	SUM
Model with triggers	- 4.97	- 0.75	- 1.6	- 0.2	- 7.52	0	- 0.53	0	- 0.17	- 0.7	0	- 0.2	- 0.91	- 2.31	- 3.42	- 11.64
Markov Model	- 6.34	- 4.27	- 2.48	- 2.19	- 15.28	- 1.23	- 0.53	- 0.53	- 2.07	- 4.36	- 1.23	- 2.19	- 2.14	- 7.14	-12.7	- 32.42

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

Other Work…

Anind Dey

Julian Ramos

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.

© 2020 Carnegie Mellon University

[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution.