Artificial Intelligence and Robotics at Oregon State University

Tom Dietterich, Distinguished Professor

24 Faculty + 3 Pending

Artificial Intelligence

- Machine Learning
- Planning and Control
- Signal Processing
- Computer Vision
- Natural Language Processing
- Human/AI Collaboration
- Genomics and Medicine
- Ecology
- Power Grid Management

Robotics

- Locomotion
  - Walking
  - Running
  - Snakes
  - Spiders
- Underwater
- Aerial
- Ground
- Prosthetics
- Human/Robot Collaboration

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Machine Learning and Data Mining

- Anomaly Detection
- Robust Machine Learning
- Causal Models for Science

Tom Dietterich

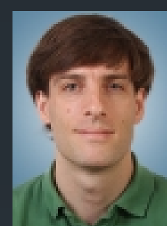

Alan Fern

Rebecca Hutchinson

Weng-Keen Wong

Arash Termehchy

Anomaly Detection

Alan Fern
Weng-Keen Wong
Tom Dietterich

- Given: a data set of tuples $x_1, \dots, x_N$
  - most of the x_i are "normal"
  - small fraction (e.g., 0.01 → 0.0001) are anomalies
- Find: the anomalies
- Applications
  - Cybersecurity
  - Data cleaning
  - Robust machine learning

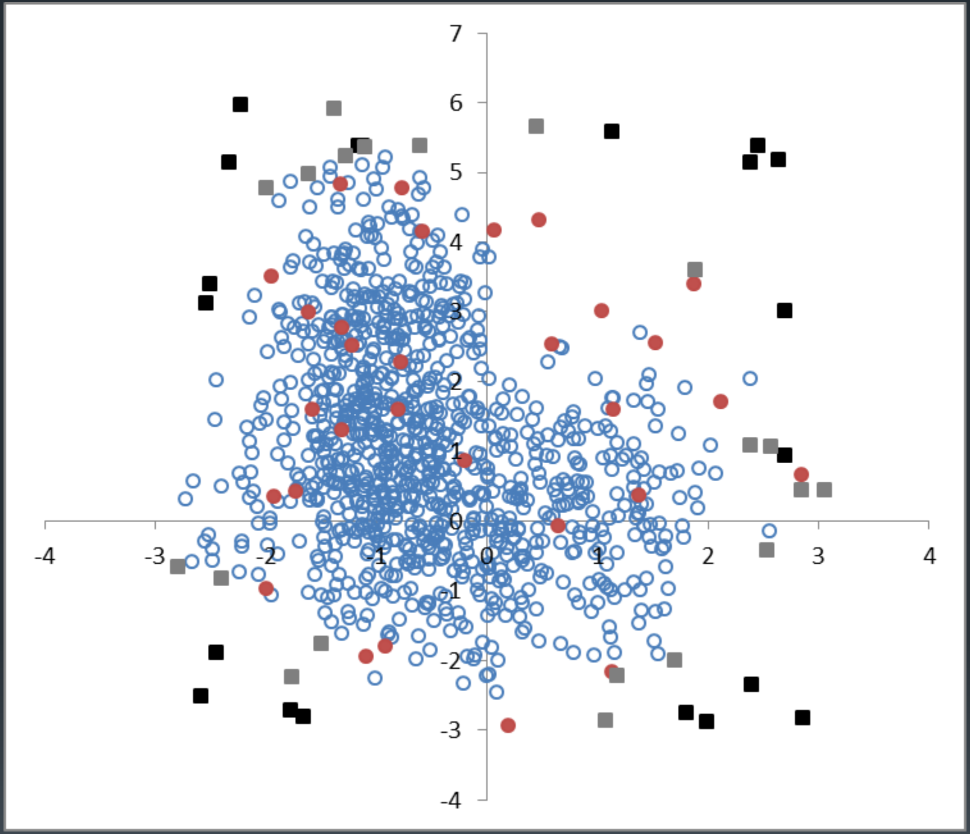

Benchmarking Study

Tom Dietterich
Andrew Emmott

- 24,800 benchmark data sets
  - constructed by manipulating existing supervised learning data sets to control
    - fraction of anomalies
    - difficulty of each point
    - clusteredness of the anomalies
    - number of irrelevant features

Oregon State UNIVERSITY

Anomaly Detection Algorithms

- Density Estimation
  - Ensemble of Gaussian Mixture Models (EGMM)
  - Robust Kernel Density Estimation (RKDE)
- Quantile Methods
  - One-class Support Vector Machines (OCSVM)
  - Support Vector Data Description (SVDD)
- Distance/Isolation Methods
  - Isolation Forest (IFOR, SCIF)
  - Local Outlier Factor (LOF)

AUC Results:

Effect of algorithm while controlling for all other factors

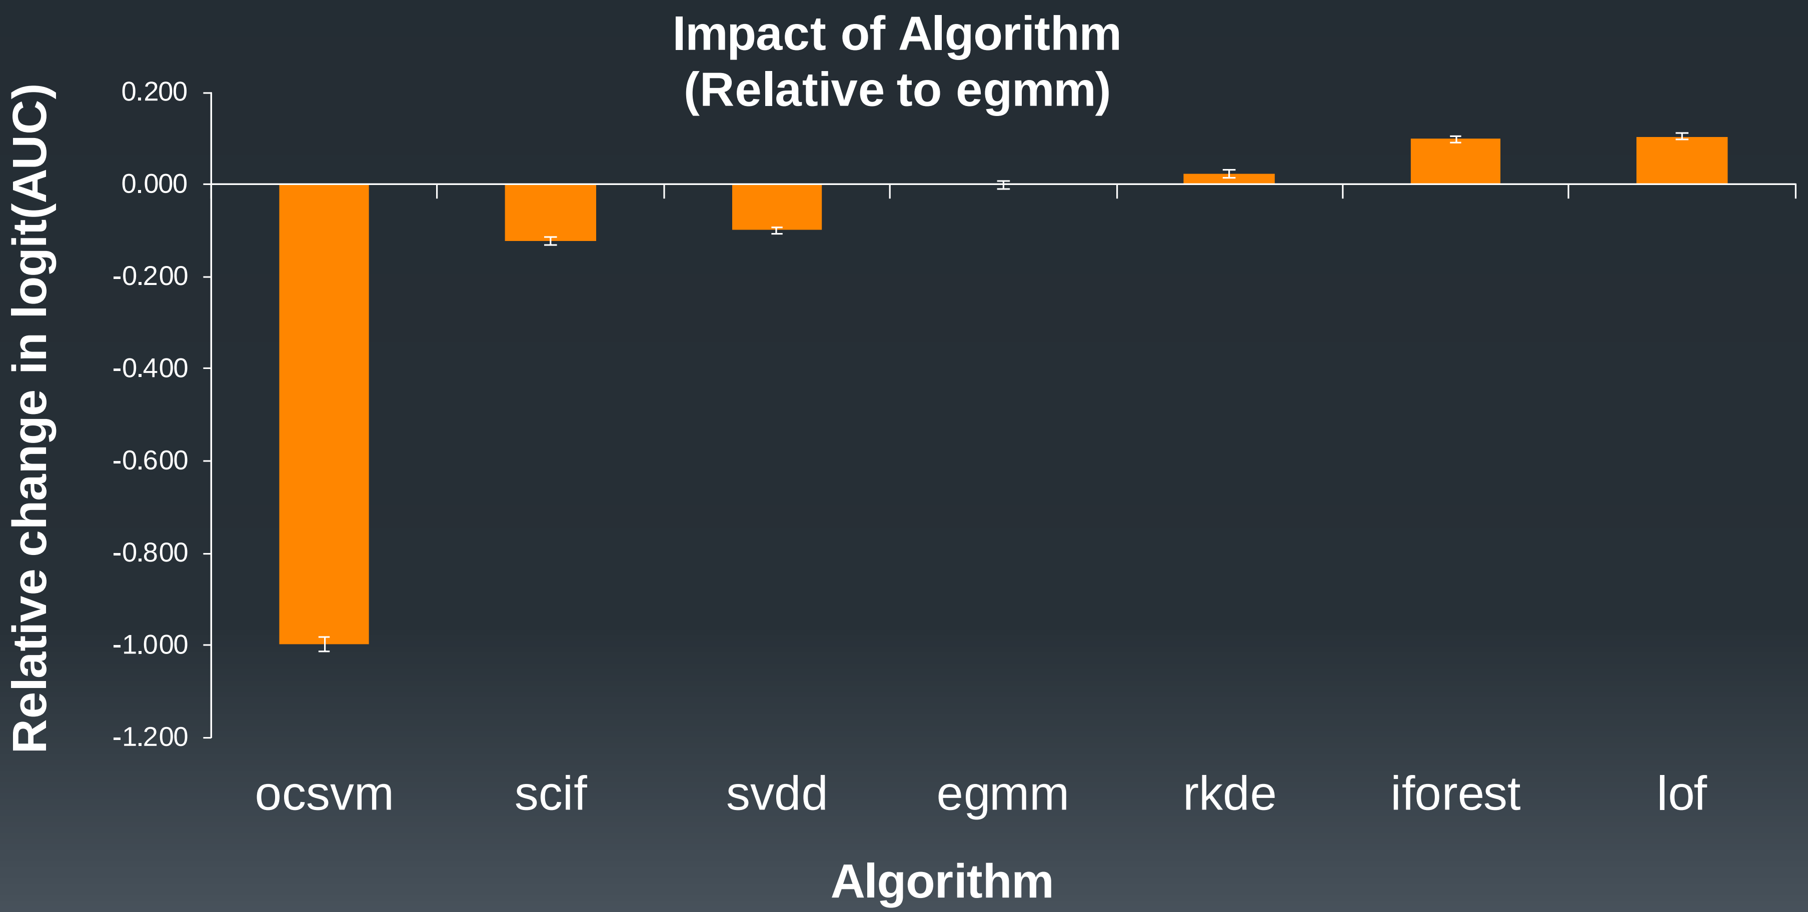

Anomaly Detection Summary

- Isolation Forest (Liu, Ting, Zhou, 2007, 2012) works very well
  - Accurate
  - Fast
  - Scalable

- Next Steps:
  - Rerunning this study to include more algorithms
    - Angle-Based Outlier Detector (ABOD)
    - Lightweight Online Detector of Anomalies (LODA)
  - PAC Theory of Anomaly Detection
    - Siddiqui, Fern, & Dietterich (UAI 2016)
  - Explanations of Anomalies
    - Siddiqui, Fern, Dietterich, Das (ODD 2015)

Applications of Anomaly Detection

- Cybersecurity
  - Detecting cyber attacks
  - Detecting insider threats
- Data Cleaning
- Robust Machine Learning

Data Cleaning for the TAHMO Project

- Africa is very poorly sensed
  - Only a few dozen weather stations reliably report data to WMO (blue points in map)
  - Goal: Make Africa the best-sensed continent

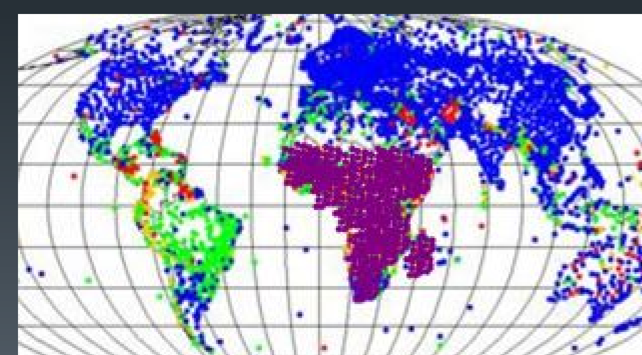

- Project TAHMO (tahmo.org)
  - TU-DELFT & Oregon State University
  - Design low-cost weather station
  - Deploy 20,000 such stations across Africa
  - Create data products (e.g., drought assessments, inundation estimates)

Data Quality Control

- An ideal method should produce two things given raw data:

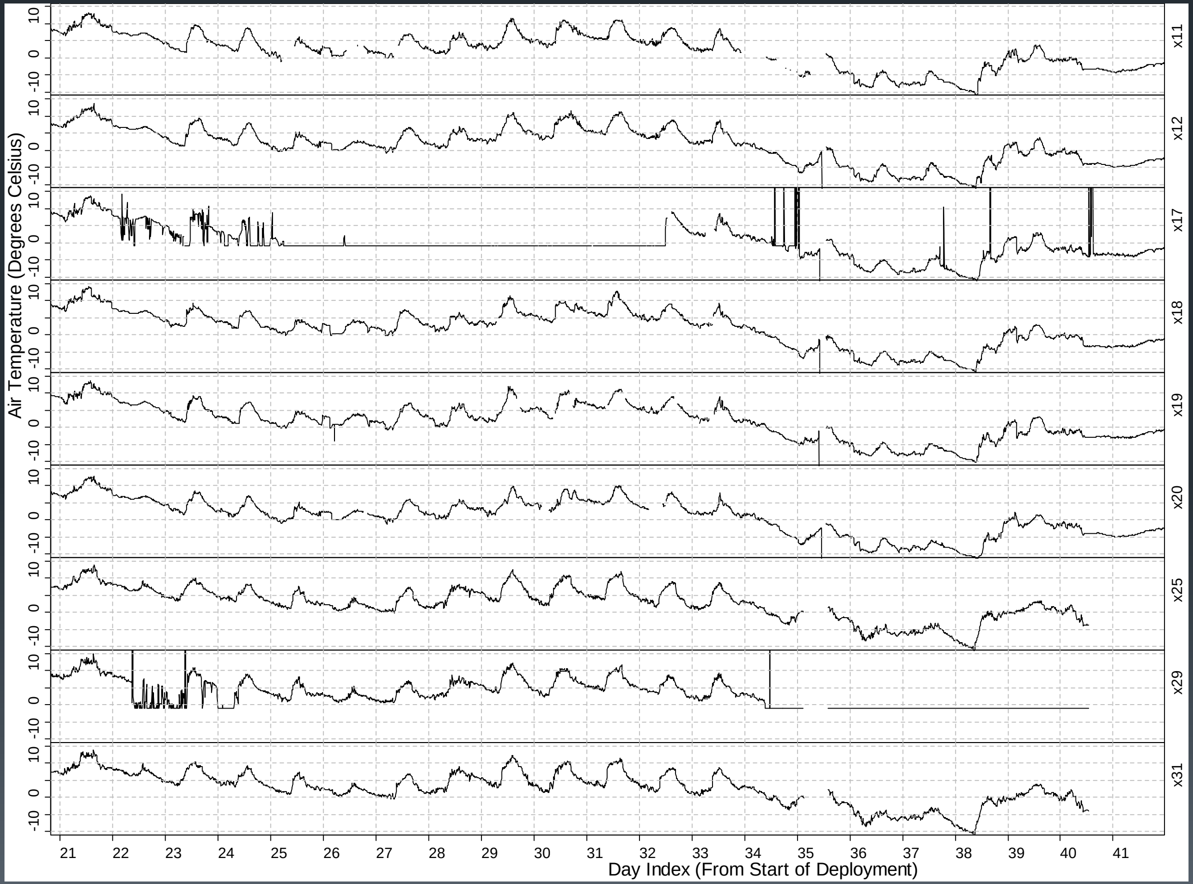

Data Quality Control

- An ideal method should produce two things given raw data:
  - A label that marks anomalies

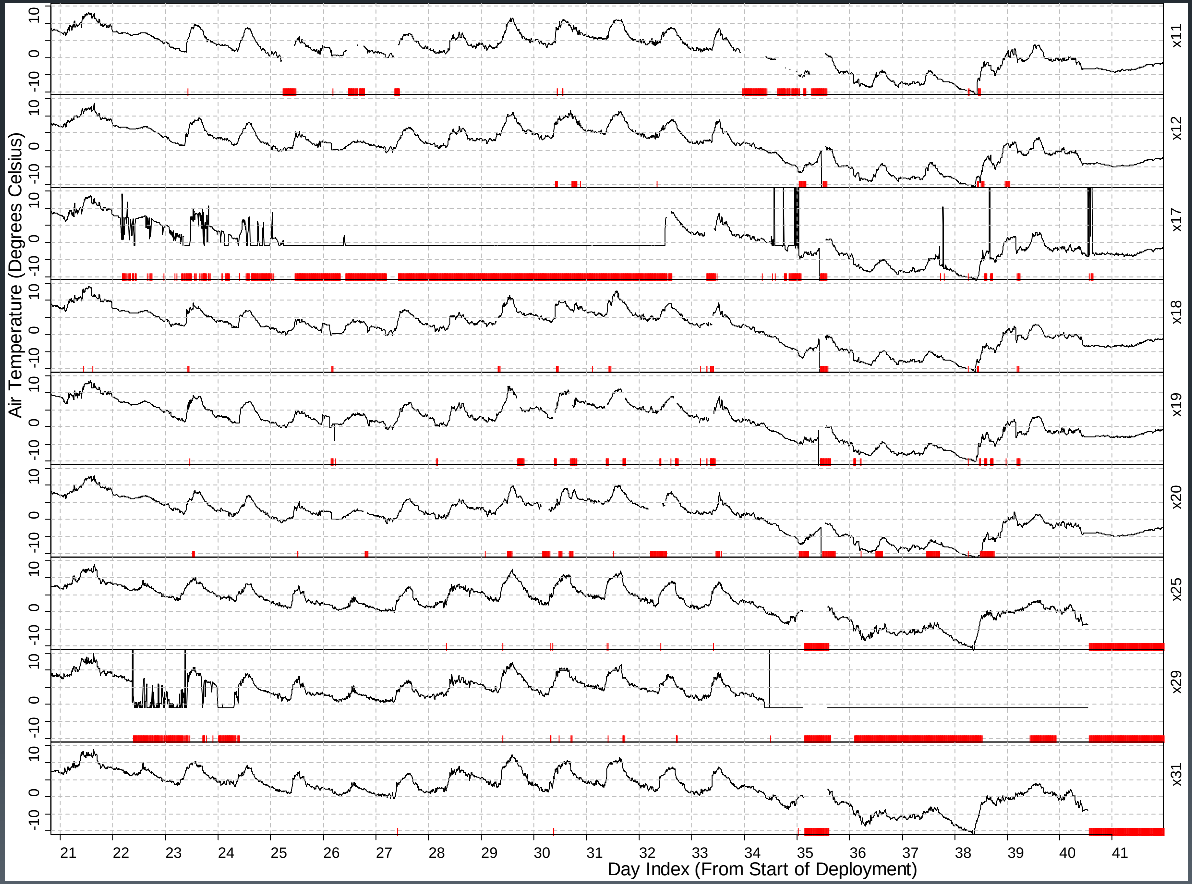

Data Quality Control

- An ideal method should produce two things given raw data:
  - A label that marks anomalies
  - An imputation of the true value (with some confidence measure)

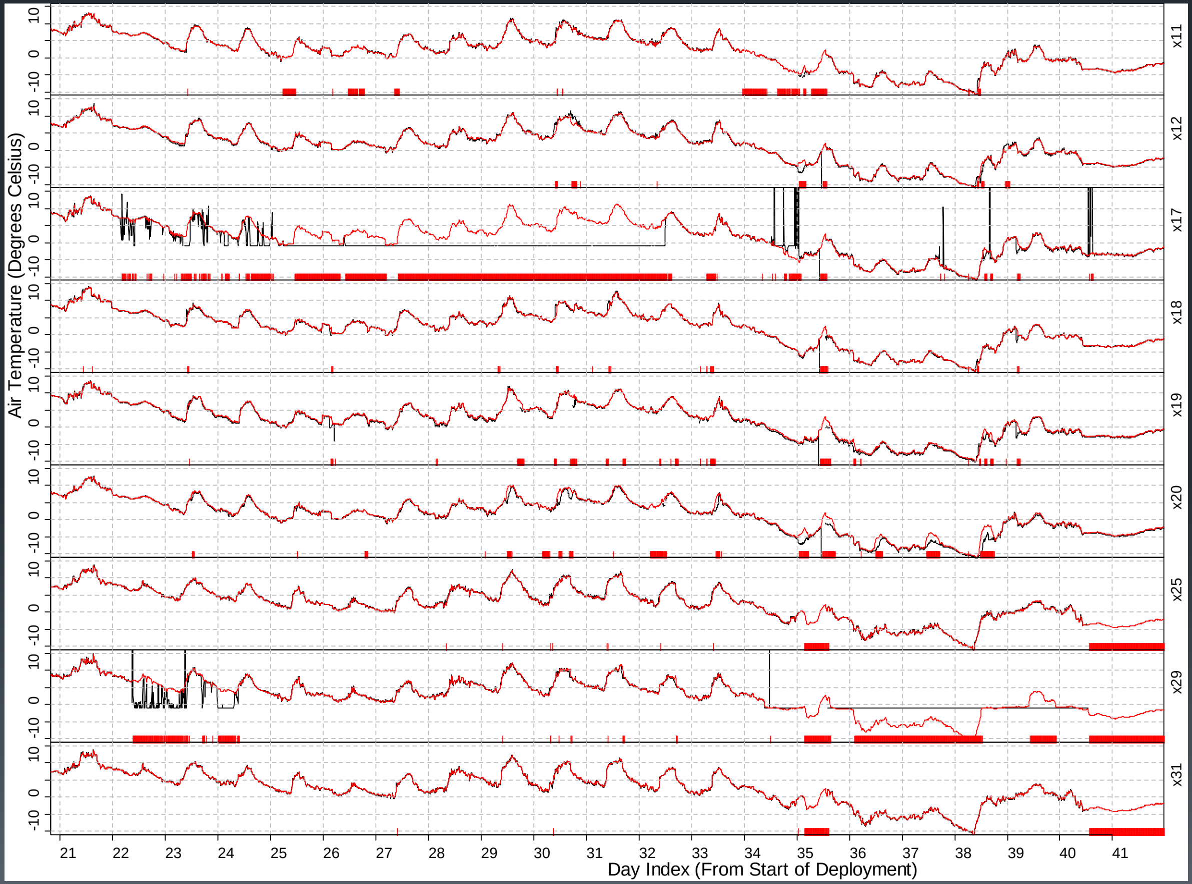

Dereszynski &, Dietterich, ACM TOS 2011.

SENSOR-DX Architecture

- General architecture for Internet of Things QC
- Find the most likely explanation of the observed anomaly scores

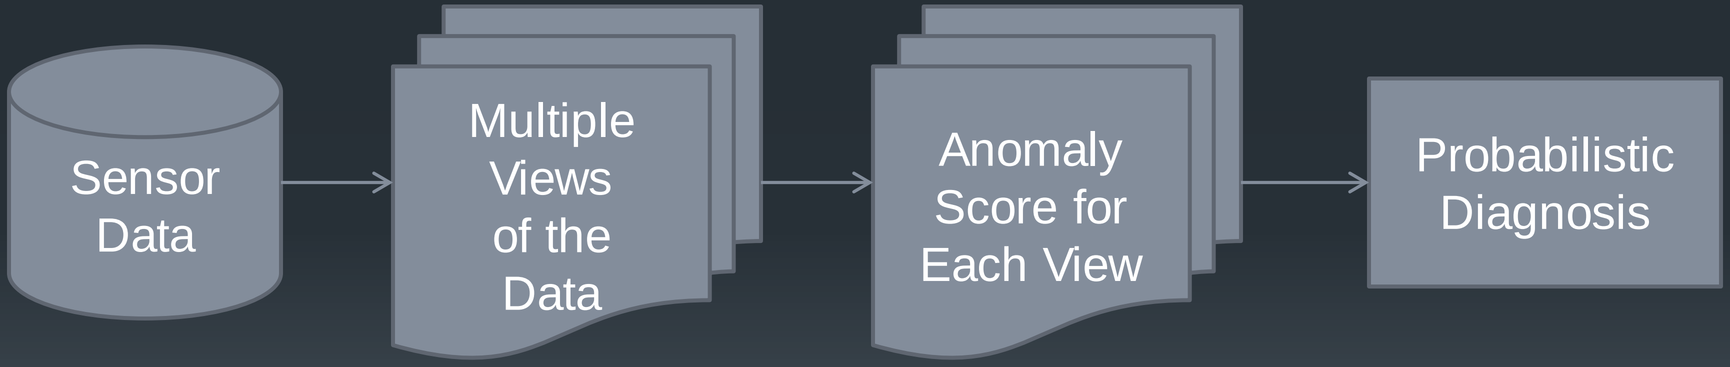

A Very Simple Example

- Weather station with two thermometers

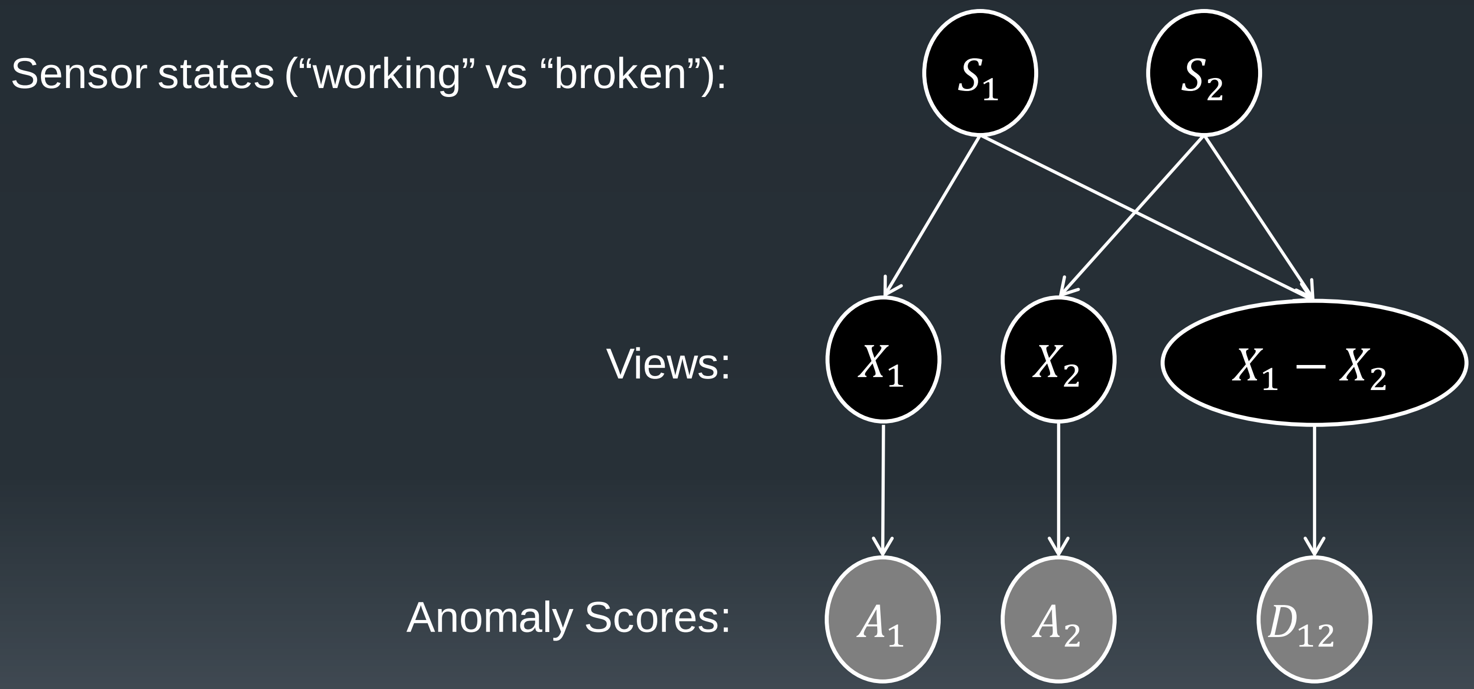

Example Result

Using $P(A_1|s_1)$ and $P(D_{12}|s_1, s_2)$

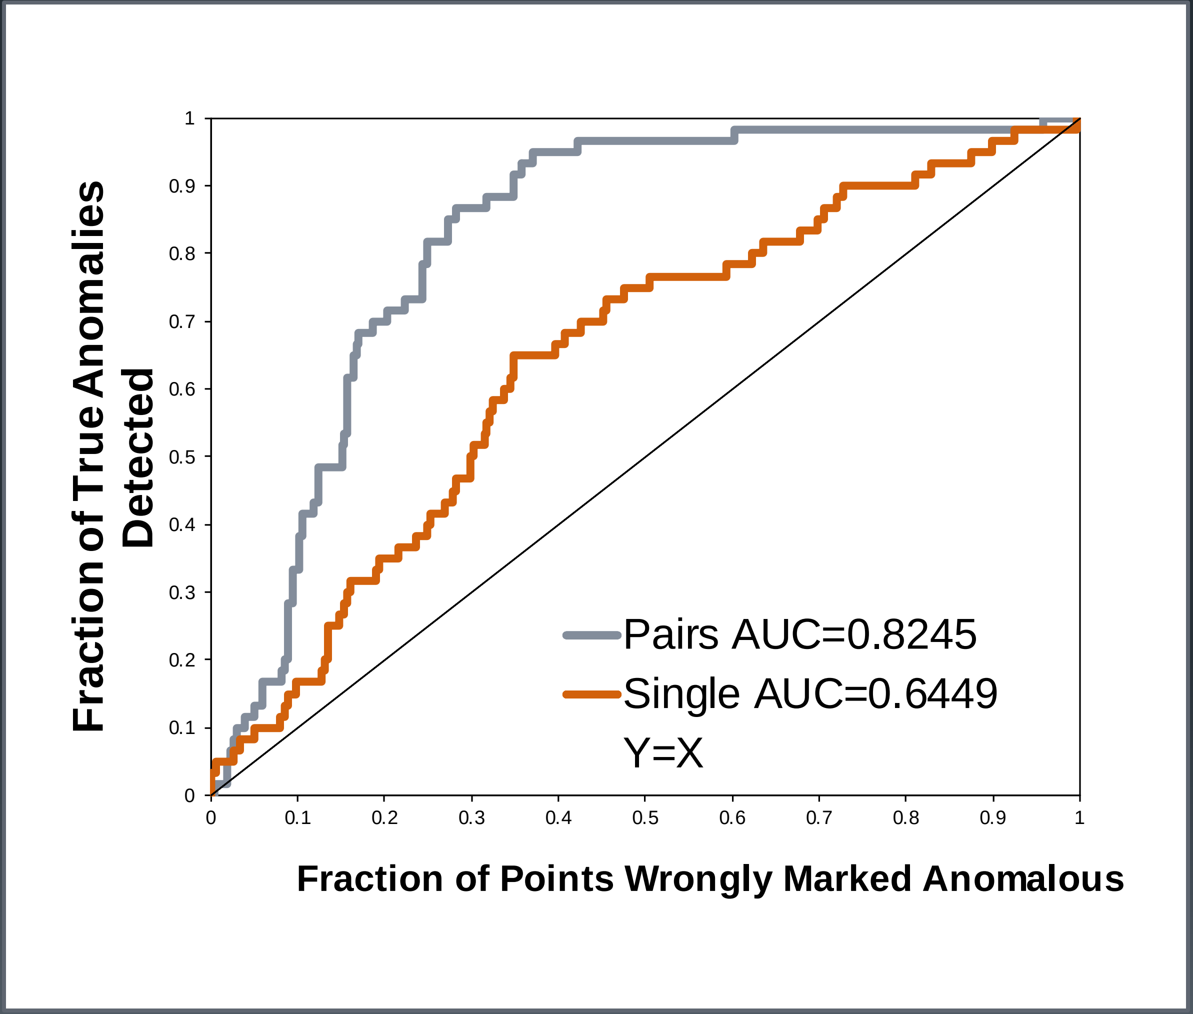

Robust Machine Learning

- Open Category Learning
  - test data contain data points belonging to "new" classes
- Non-Stationary Learning
  - test data contain data from a very different distribution

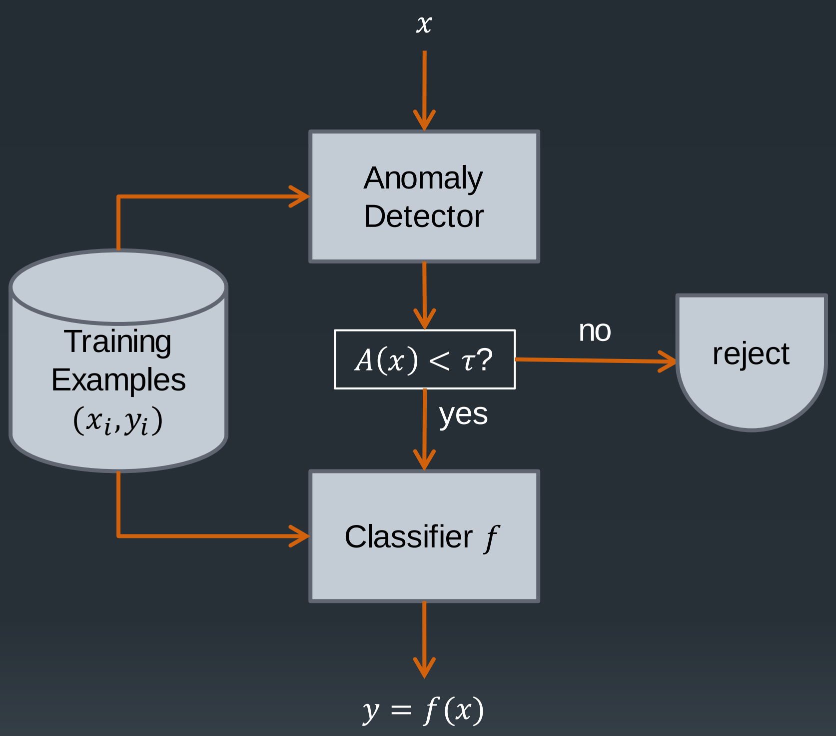

Fern, Dietterich, Juozapaitis, unpublished

Causal Models for Science

- Modeling Bird Migration
  - Input data: “Citizen Science” observations
  - Output model: predicts when and where birds will fly

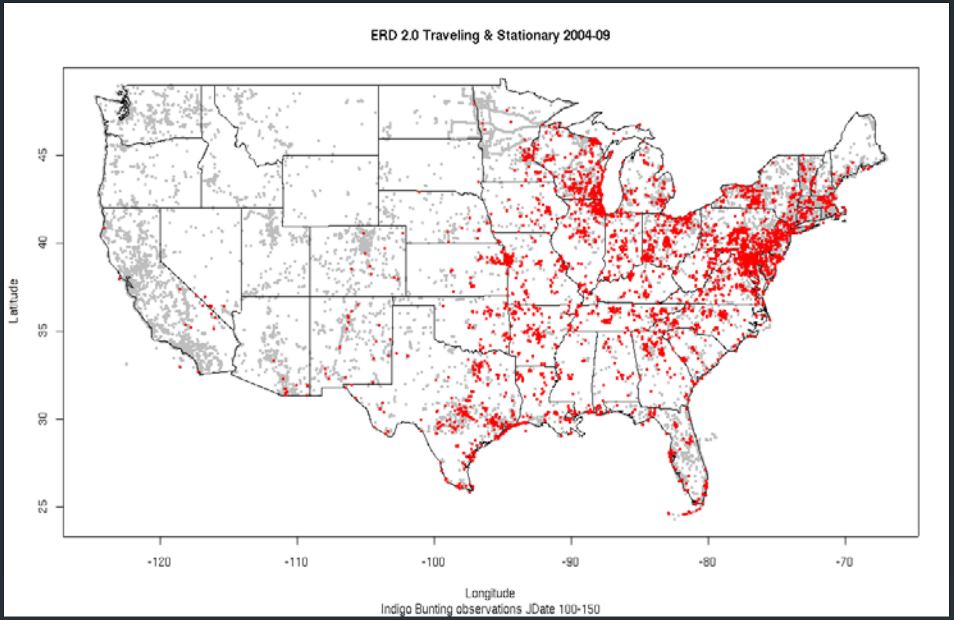

Sheldon, Sun, Liu, Dietterich, Cornell Lab of Ornithology

Hidden Markov Model

$P(x_1)$: Initial State Distribution
$P(x_t|x_{t-1})$: State transition function

- Individual Model
  - Model of a single bird's location over time
- Population Model
  - Model of the spatial distribution of a population of birds over time

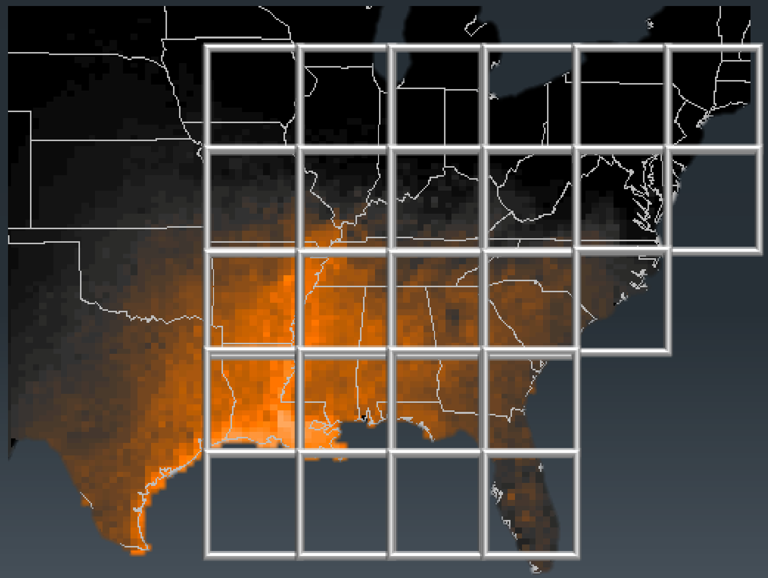

Fitting the model
The data we wish we had:

- Tracks of individual birds over time
- Weather at every location

macworld.com

www.azoresbioportal.angra.uac.pt

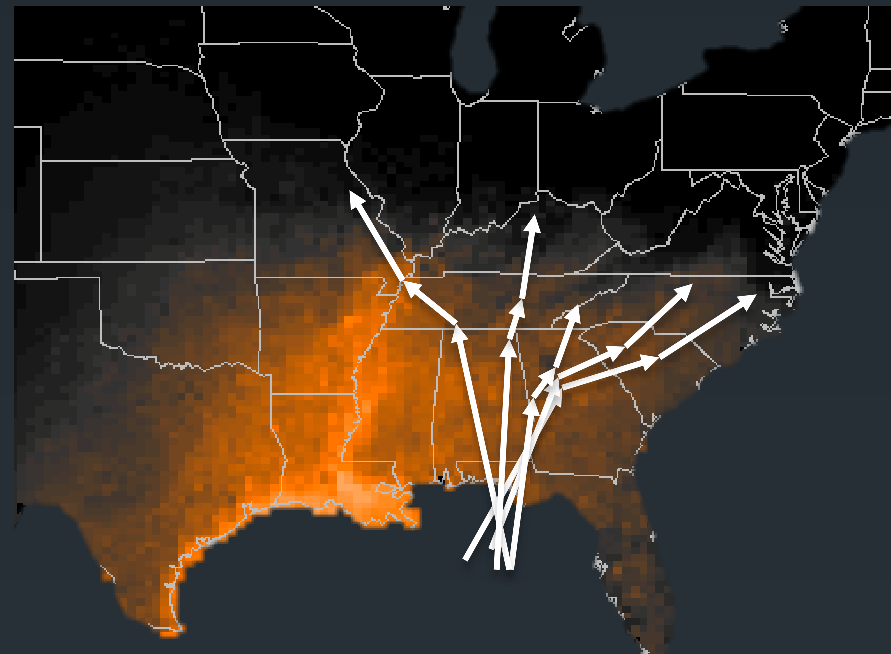

This would give us points $(x_{t,t+1}(i), x_t(i), x_{t+1}(j))$ to which we could fit our model

Challenge: Aggregate anonymous counts

- We do not observe the behavior of individual birds: $x_t(i)$
- We only obtain information about aggregated counts of birds: $n_t(i)$

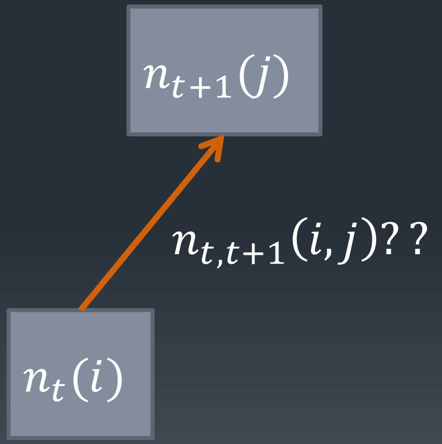

Solution: Collective Graphical Models

Individual model:
Markov chain on grid cells

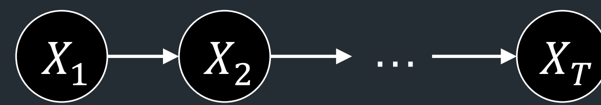

Population model:
M iid copies of individual model

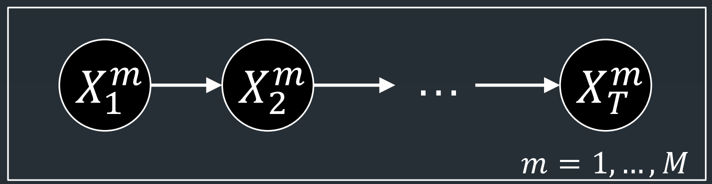

Derive aggregate counts

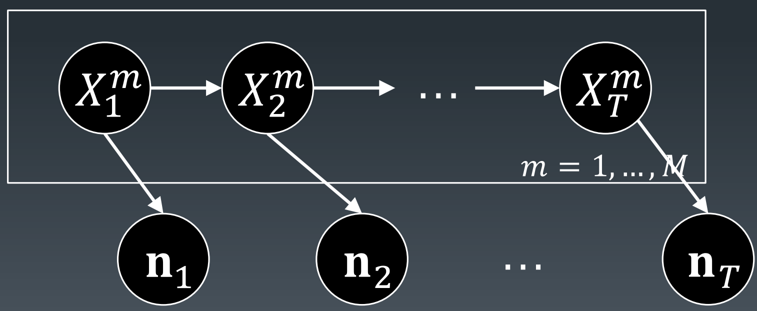

Solution:
Collective Graphical Models (2)

Derive aggregate counts

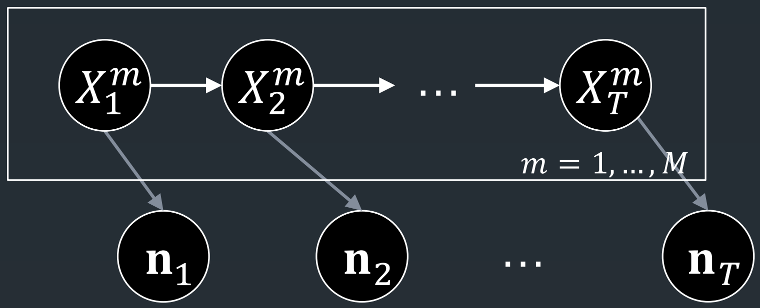

Attach Noisy Observations

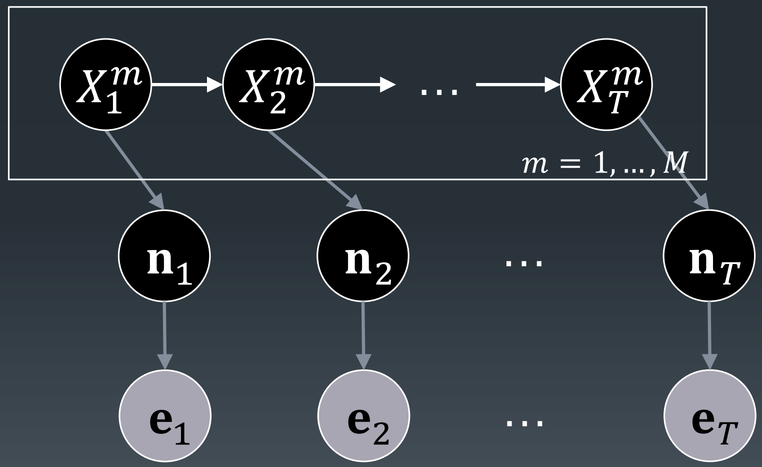

Solution:
Collective Graphical Models (3)

Attach Noisy Observations

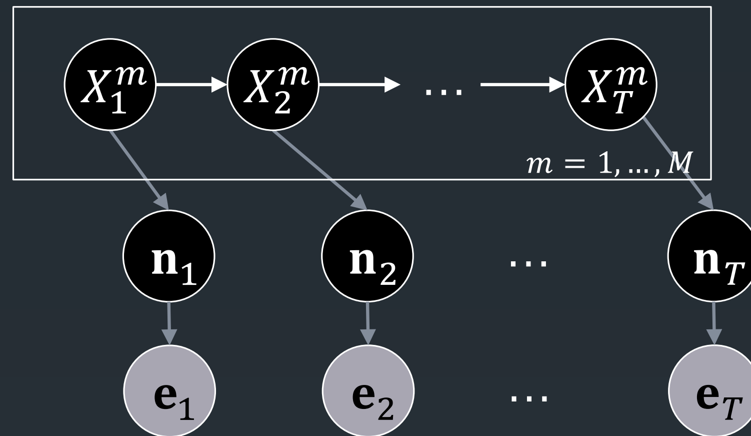

Marginalize out individuals: chain-structured model on sufficient statistics

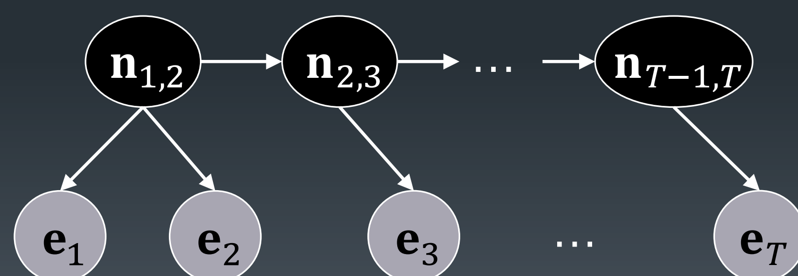

Transition counts

MAP Inference in the Collective Graphical Model

- Generic Message Passing
- Non-linear belief propagation
  - NLBP-NAÏVE
    - same message-passing schedule as in standard BP
  - NLBP-FEAS
    - special schedule
    - dampened updates
    - guarantees feasibility at each step

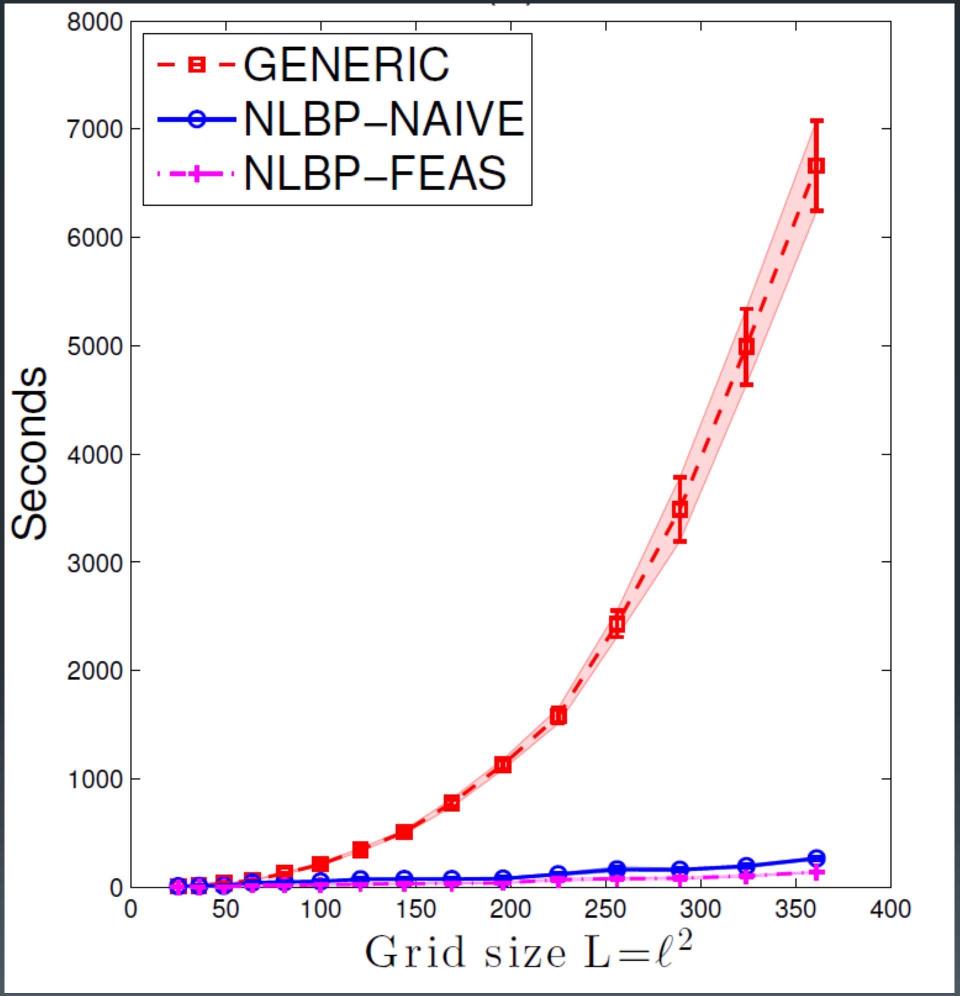

Eastern Wood Peewee

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Planning and Control

- Experiment Planning for Microbial Fuel Cells
- Preventing or Reducing Blackouts in the Power Grid
- Managing Endangered Species

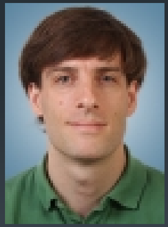
Alan Fern


Xiaoli Fern


Tom Dietterich


Prasad Tadepalli

- Given: a simulator of the system
- Find: find near-optimal ways to control the system
- Methods:
  - Monte Carlo Tree Search
  - Reinforcement Learning
  - Bayesian Optimization

Parallel Bayesian Optimization

Xiaoli Fern

Previous Experiments

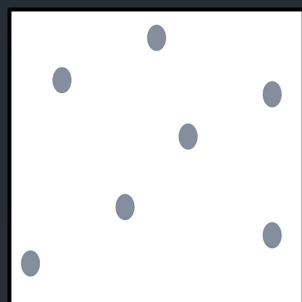

Posterior Model

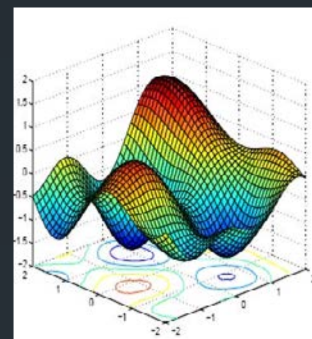

Choose Next K Experiments

Run K Experiments

How many to run in parallel?
When to start new experiments?
Which experiments to start?

Simulation-Based Electric Grid Stabilization

Alan Fern

Northeast Blackout, 2003

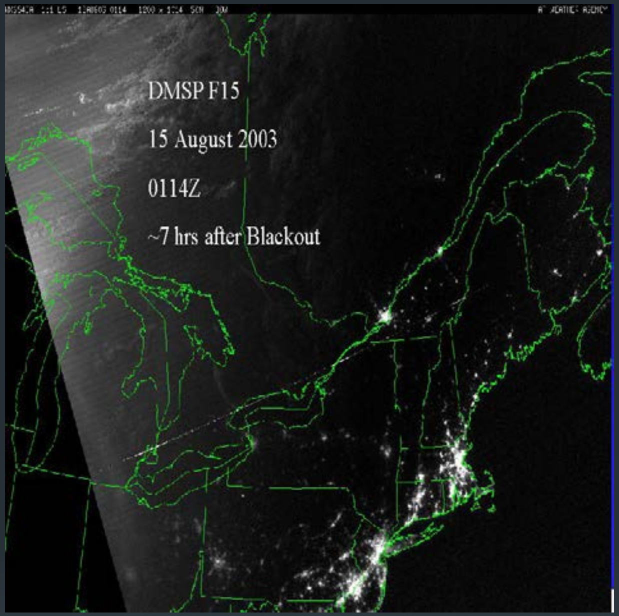

Credit: Air Force Weather Agency (AFWA)

Credit: Jonathan Fickies/Getty Images

- How can AI help?

Electrical Grid Stabilization

- *N-1* security

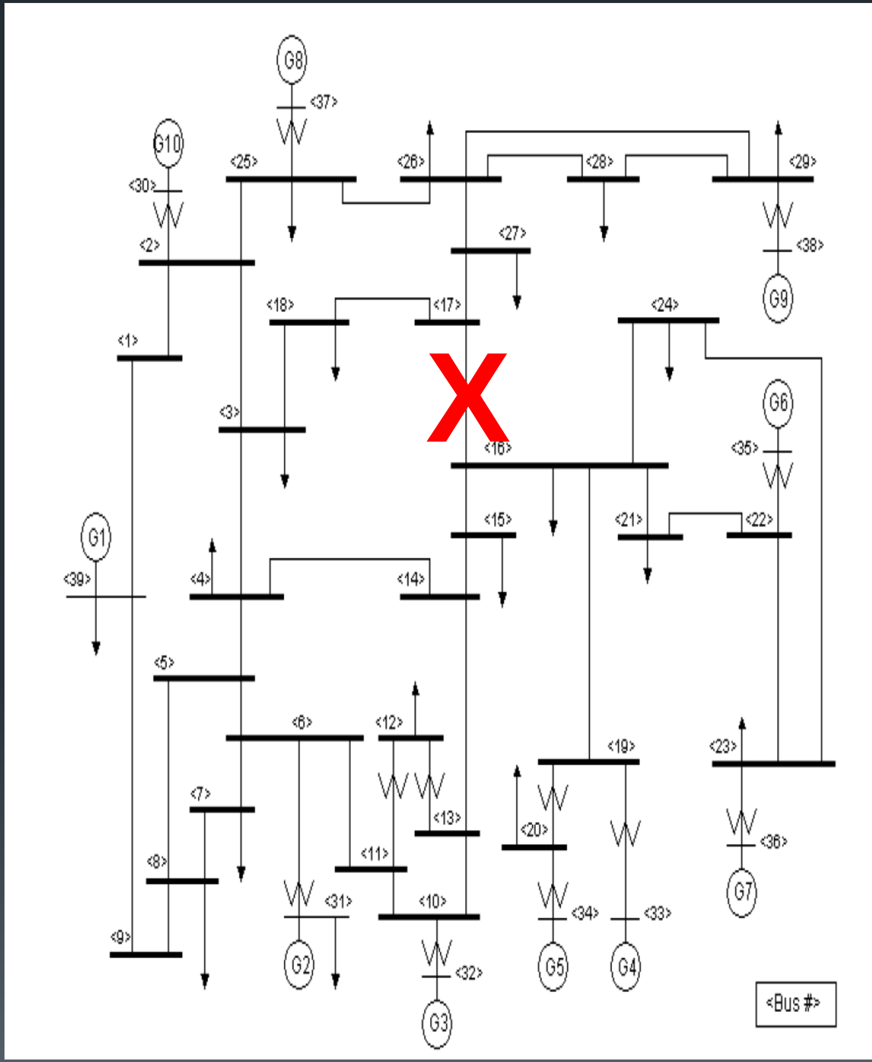

Electrical Grid Stabilization

- What about *N-2* ?

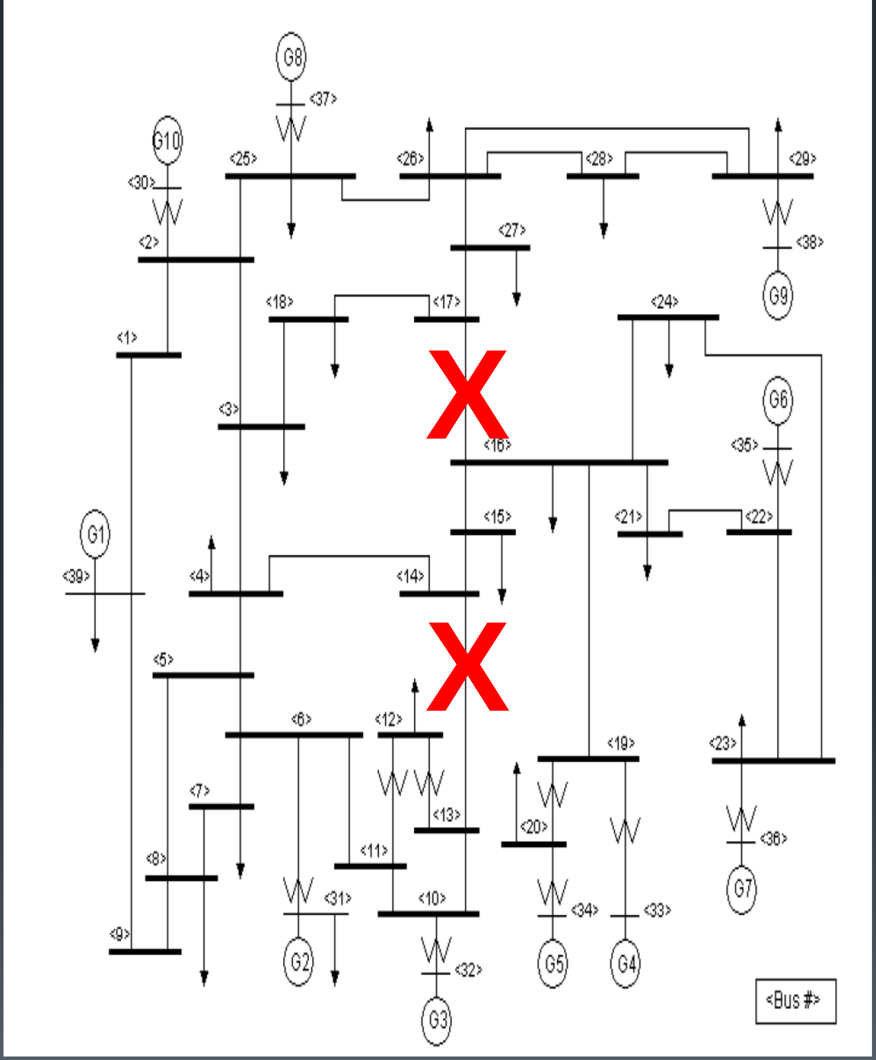

Electrical Grid Stabilization

- Getting hot…

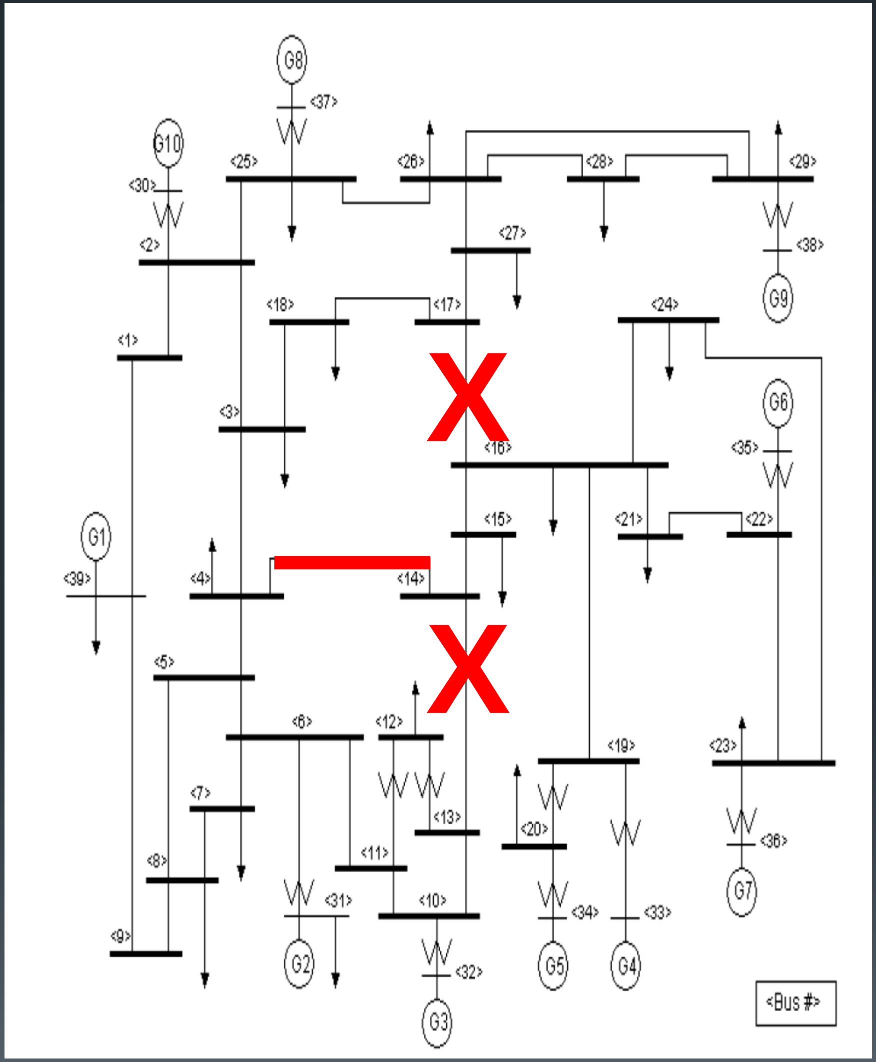

Electrical Grid Stabilization

- Avoid Cascading failure
- Prevent or minimize damage by taking appropriate **remedial action schemes**
  - Load shedding
  - Islanding
- Island or Load Shed?
  - Where?
  - When?

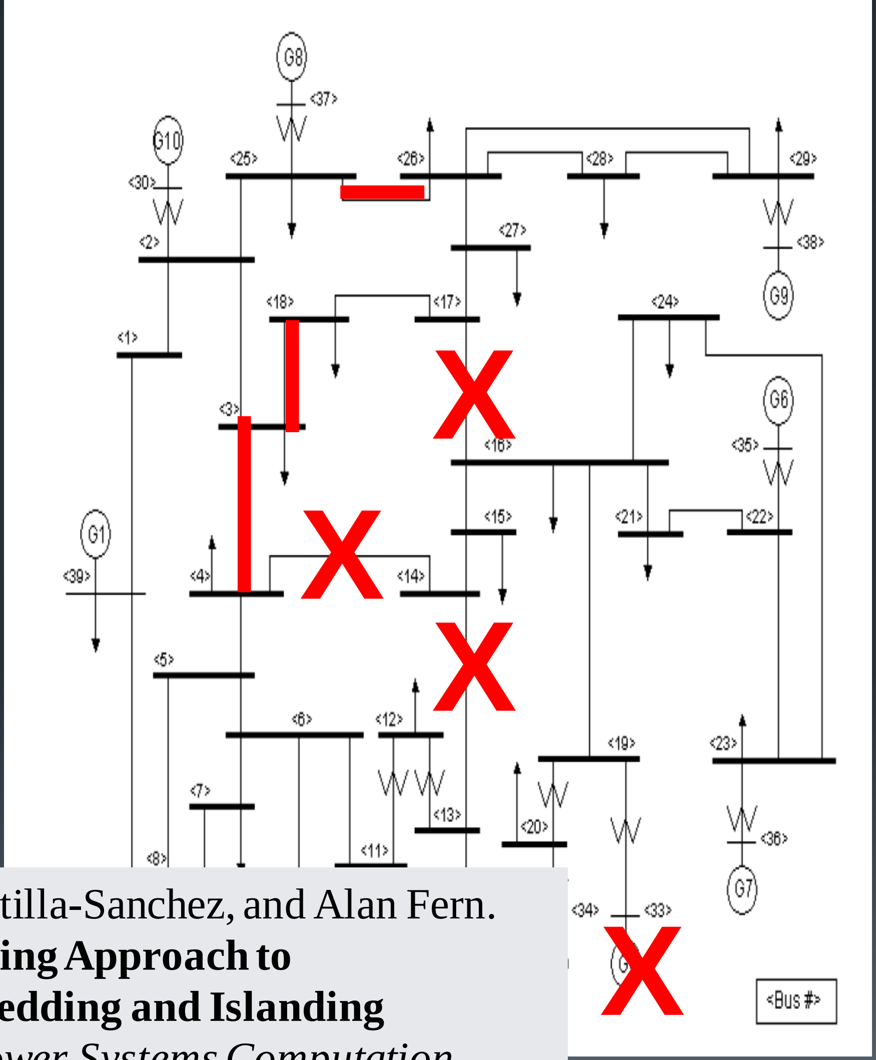

Rich Meier, Eduardo Cotilla-Sanchez, and Alan Fern. (2014). **A Policy Switching Approach to Consolidating Load Shedding and Islanding Protection Schemes.** *Power Systems Computation Conference*

OSU
Oregon State
UNIVERSITY

Conservation Planning

Alan Fern

Red area: current location of the Red-Cockaded Woodpecker

Goal: purchase and conserve land parcels to maximize population spread of target species

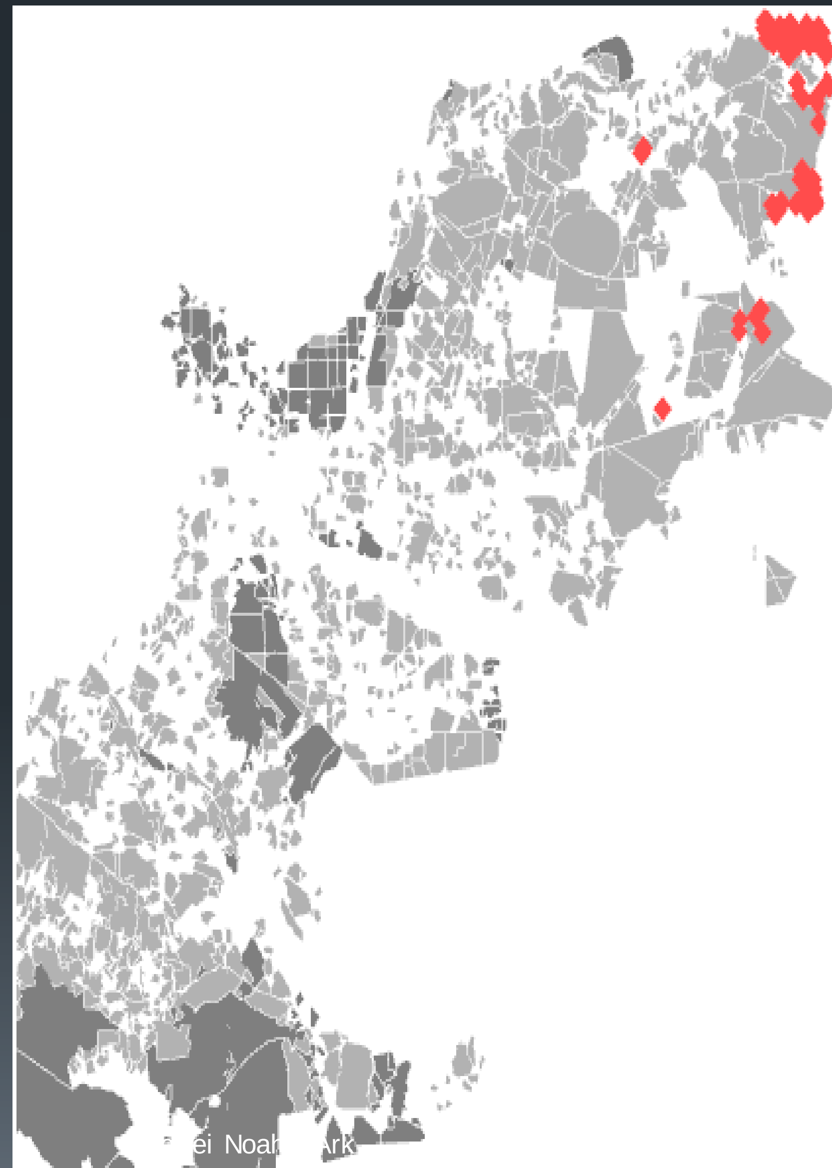

Red-Cockaded Woodpecker

Conservation Planning

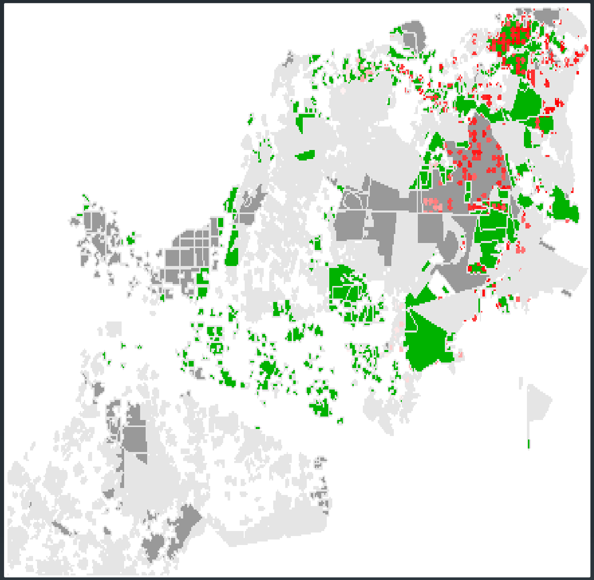

Year 1

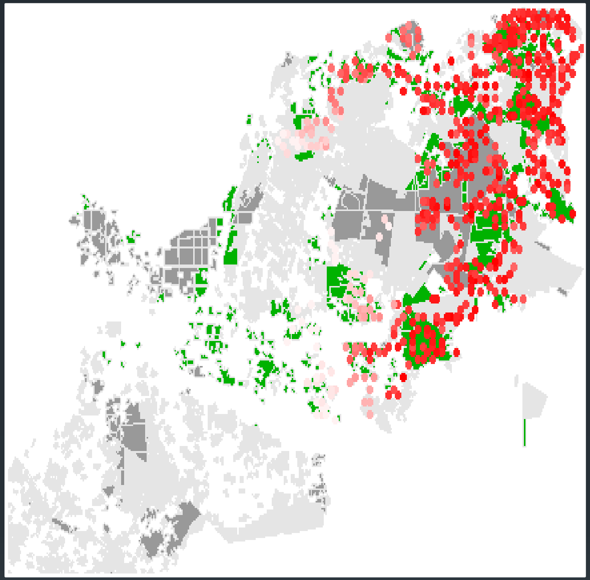

Year 5

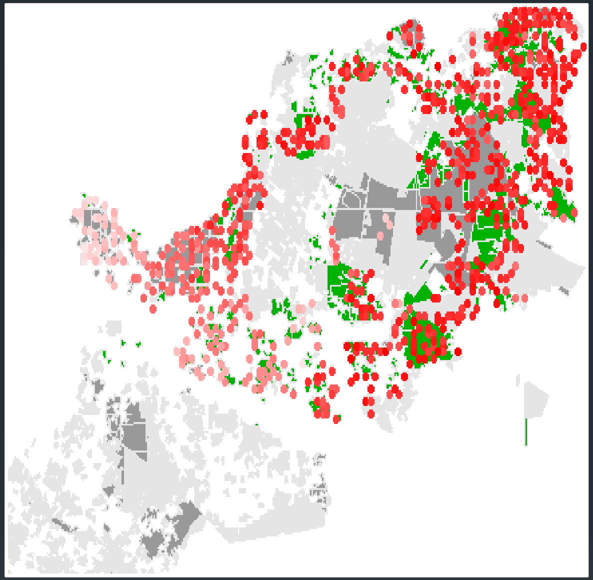

Year 10

Being implemented by the The Conservation Fund

The Big Picture

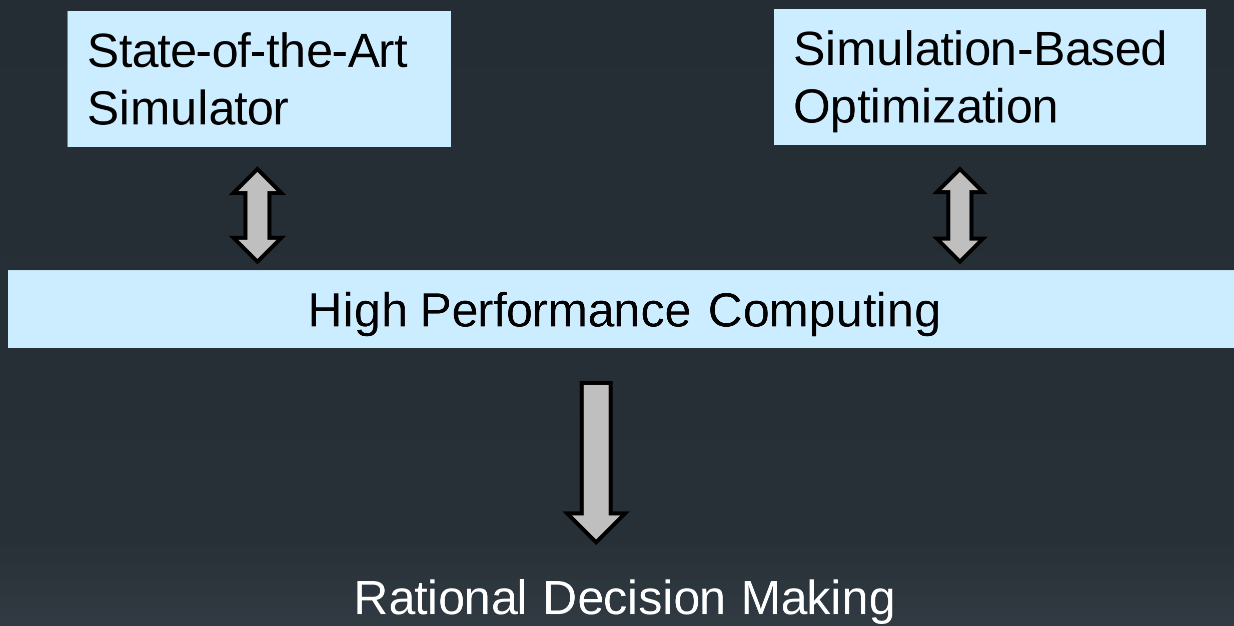

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Computer Vision and Acoustics

- Video Action Recognition (Sinisa)
- Object Recognition and Segmentation (Fuxin, Sinisa)
- Video Object Tracking (Fuxin, Sinisa, Alan)
- Bird Song Recognition (Xiaoli, Raviv)


Sinisa Todorovic


Fuxin Li

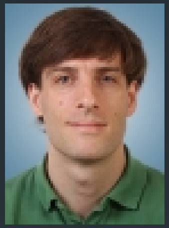
Alan Fern

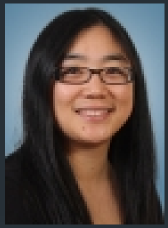
Xiaoli Fern


Raviv Raich

Action Recognition in the Wild

Behrooz Mahasseni
Sinisa Todorovic
CVPR 2016

- No priors on camera viewpoint
- Large number of action classes
- Different motion patterns
- Background Clutter and Occlusion

Fishing

Running

Wind Surfing

Oregon State University

State of the art

- Deep Learning Approaches
  - DCNN + (Time Pooling / LSTM / Graphical Models)
- Current Trend
  - Deeper Networks + More training data → Higher accuracy
- Our Idea
  - Augment with complementary data **ONLY in TRAINING**
    - Dynamics of the motion pattern
    - 3D human skeleton sequences

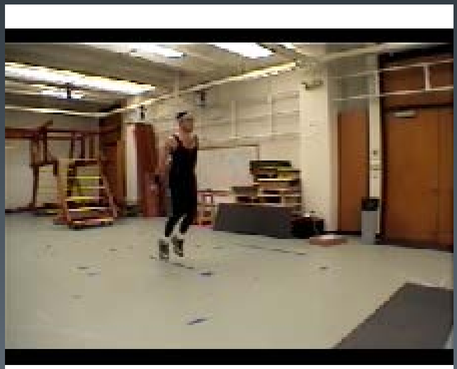

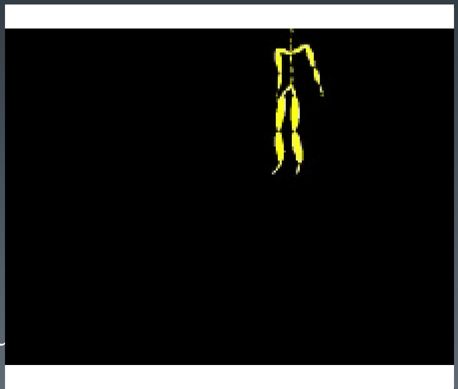

Oregon State UNIVERSITY

Hu

Result

- Dataset: Sport1M

Method	Hit@1	Hit@5
Single Frame	59.3	77.7
LSTM	71.3	89.9
[1]	60.9	80.2
[2]	72.1	90.6
[3]	61.1	85.2
R-LSTM	**75.9**	**91.7**

[1] Karpathy et al. Large-scale video classification with convolutional neural networks. In CVPR, 2014
[2] Ng et al. Beyond short snippets: Deep networks for video classification, arXiv2015
[3] Tran et al. C3D: generic features for video analysis. CoRR 2014

Result

- Dataset: UCF101, HMDB-51

Method	UCF101	HMDB-51
[1]	65.4	-
[2]	75.8	44.1
[3]	71.12	-
[4]	72.8	40.5
[5]	79.34	-
[6]	85.2	-
R-LSTM	**86.9**	**55.3**

[1] Karpathy et al. Large-scale video classification with convolutional neural networks. CVPR, 2014
[2] Srivastava et al. Unsupervised learning of video representations using lstms, arXiv2015
[3] Donahue et al. Long-term recurrent convolutional networks for visual recognition and description, arXiv 2014
[4] Simonyan et al. Two-stream convolutional networks for action recognition in videos NIPS 2014
[5] Zha et al. Exploiting image-trained cnn architectures for unconstrained video classification, arXiv 2015
[6] Tran et al. C3D: generic features for video analysis, CoRR, 2014

Semantic Segmentation

Fuxin Li, et al.

- Given an image, identify the category and spatial extent of all relevant objects
- Focus on detailed object interactions

Image

Given:

Proposals

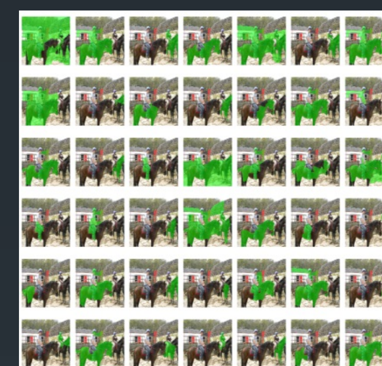

Goal:

Object Label

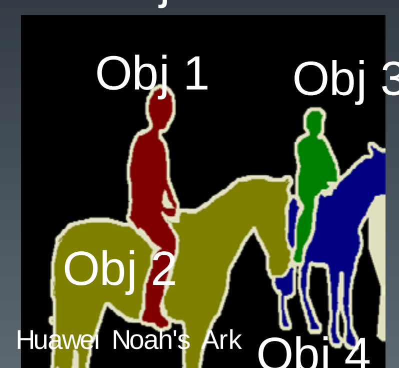

Category Label

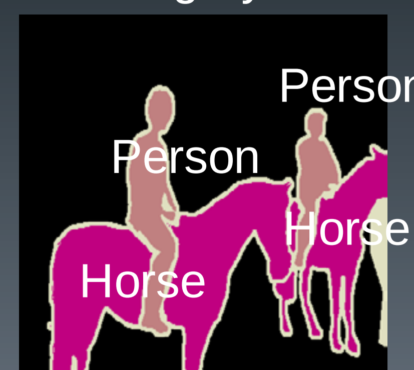

Semantic Segmentation

- Predict % overlap of categories, not objects

X: Shape/Color/Texture feature of segments

Y: Class-specific overlap

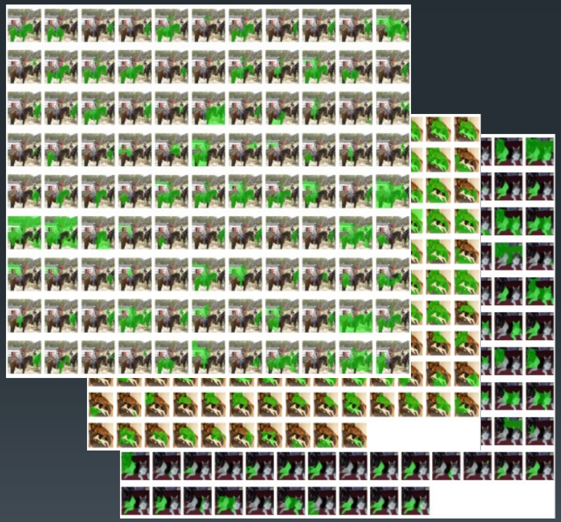

1 regressor per category

Horse 0.85

Bus 0.3

Dog 0.5, Cat 0.5

Person 0.95

Plant 0.9

Bike 0.4

Li, Carreira and Sminchisescu. CVPR 10, IJCV 12

Dissecting Segments

Image

Segment #1:
Chair 0.53
Person 0.29

Segment #2:
Chair 0.23
Person 0.36

Atomic regions:

Venn diagram

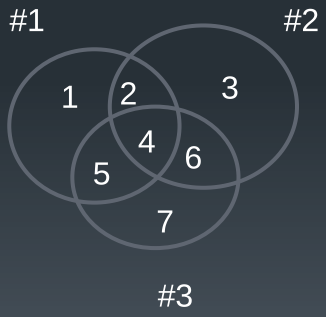

Segment #3:
Chair 0.34
Person 0.54

Segment #4:
Chair 0.19
Person 0.43

Find Most Consistent Configuration

Seg #1: Chair 0.53
Person 0.29

Seg #2: Chair 0.23
Person 0.36

Seg #3: Chair 0.34
Person 0.54

Seg #4: Chair 0.19
Person 0.43

Configuration #1

Person Chair

Configuration #2

Person Chair

Results: Winners of PASCAL VOC 2009-12

Segmentation Results: VOC2012 BETA

Competition "comp5" (train on VOC2012 data)

Average Precision (AP %)

	mean	aero plane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	dining table	dog	horse	motor bike	person	potted plant	sheep	sofa	train	tv/ monitor
O2P_SVRSEGM_CPMC_CSI [?]	**47.5**	64.0	**32.2**	45.9	34.7	46.3	59.5	**61.7**	49.4	14.8	47.9	31.2	42.5	51.3	**58.8**	**54.6**	34.9	54.6	34.7	50.6	42.2
NUS_DET_SPR_GC_SP [?]	47.3	52.9	31.0	39.8	**44.5**	**58.9**	**60.8**	52.5	49.0	**22.6**	38.1	27.5	**47.4**	52.4	46.8	51.9	**35.7**	**55.3**	**40.8**	**54.2**	**47.8**
BONN_O2PCPMC_FGT_SEGM [?]	47.0	**65.4**	29.3	**51.3**	33.4	44.2	59.8	60.3	52.5	13.6	**53.6**	**32.6**	40.3	**57.6**	57.3	49.0	33.5	53.5	29.2	47.6	37.6
BONNGC_O2P_CPMC_CSI [?]	45.4	59.3	27.9	43.9	39.8	41.4	52.2	61.5	**56.4**	13.6	44.5	26.1	42.8	51.7	57.9	51.3	29.8	45.7	28.8	49.9	43.3
BONN_CMBR_O2P_CPMC_LIN [?]	44.8	60.0	27.3	46.4	40.0	41.7	57.6	59.0	50.4	10.0	41.6	22.3	43.0	51.7	56.8	50.1	33.7	43.7	29.5	47.5	44.7
OptNBNN-CRF [?]	11.3	10.5	2.3	3.0	3.0	1.0	30.2	14.9	15.0	0.2	6.1	2.3	5.1	12.1	15.3	23.4	0.5	8.9	3.5	10.7	5.3

OSU Digital Scout

Alan Fern
Sinisa Todorovic

- Goal: Understand American football
  - Determine the start of each play
  - Track each player
  - Understand what choices each player made
  - Understand what mistakes each player made
  - Improve the play
  - Plan how to defeat this play

- Industrial collaboration with HUDL, Inc.
  - Video analytics for US high school football teams

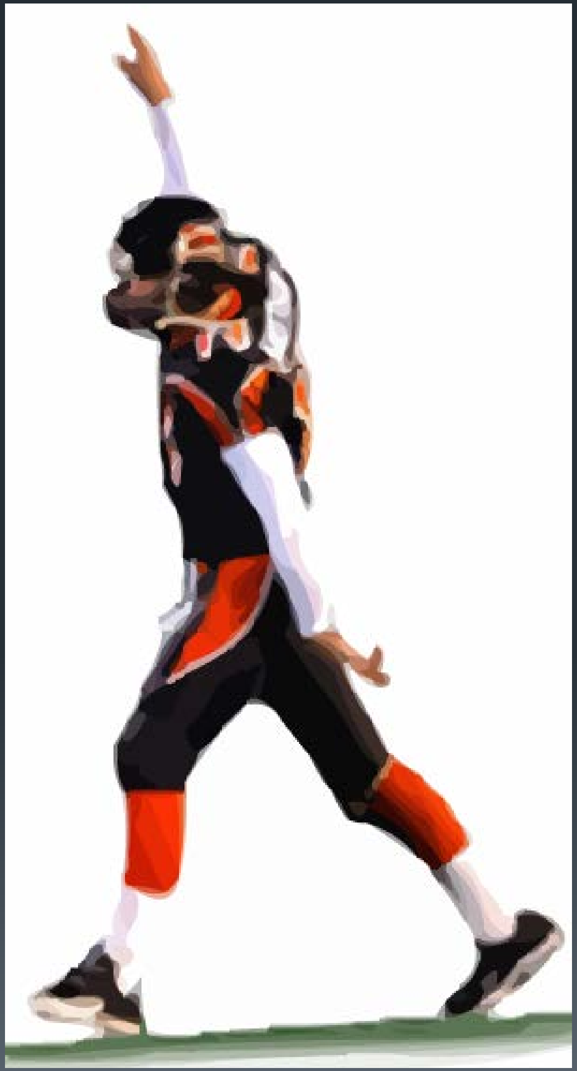

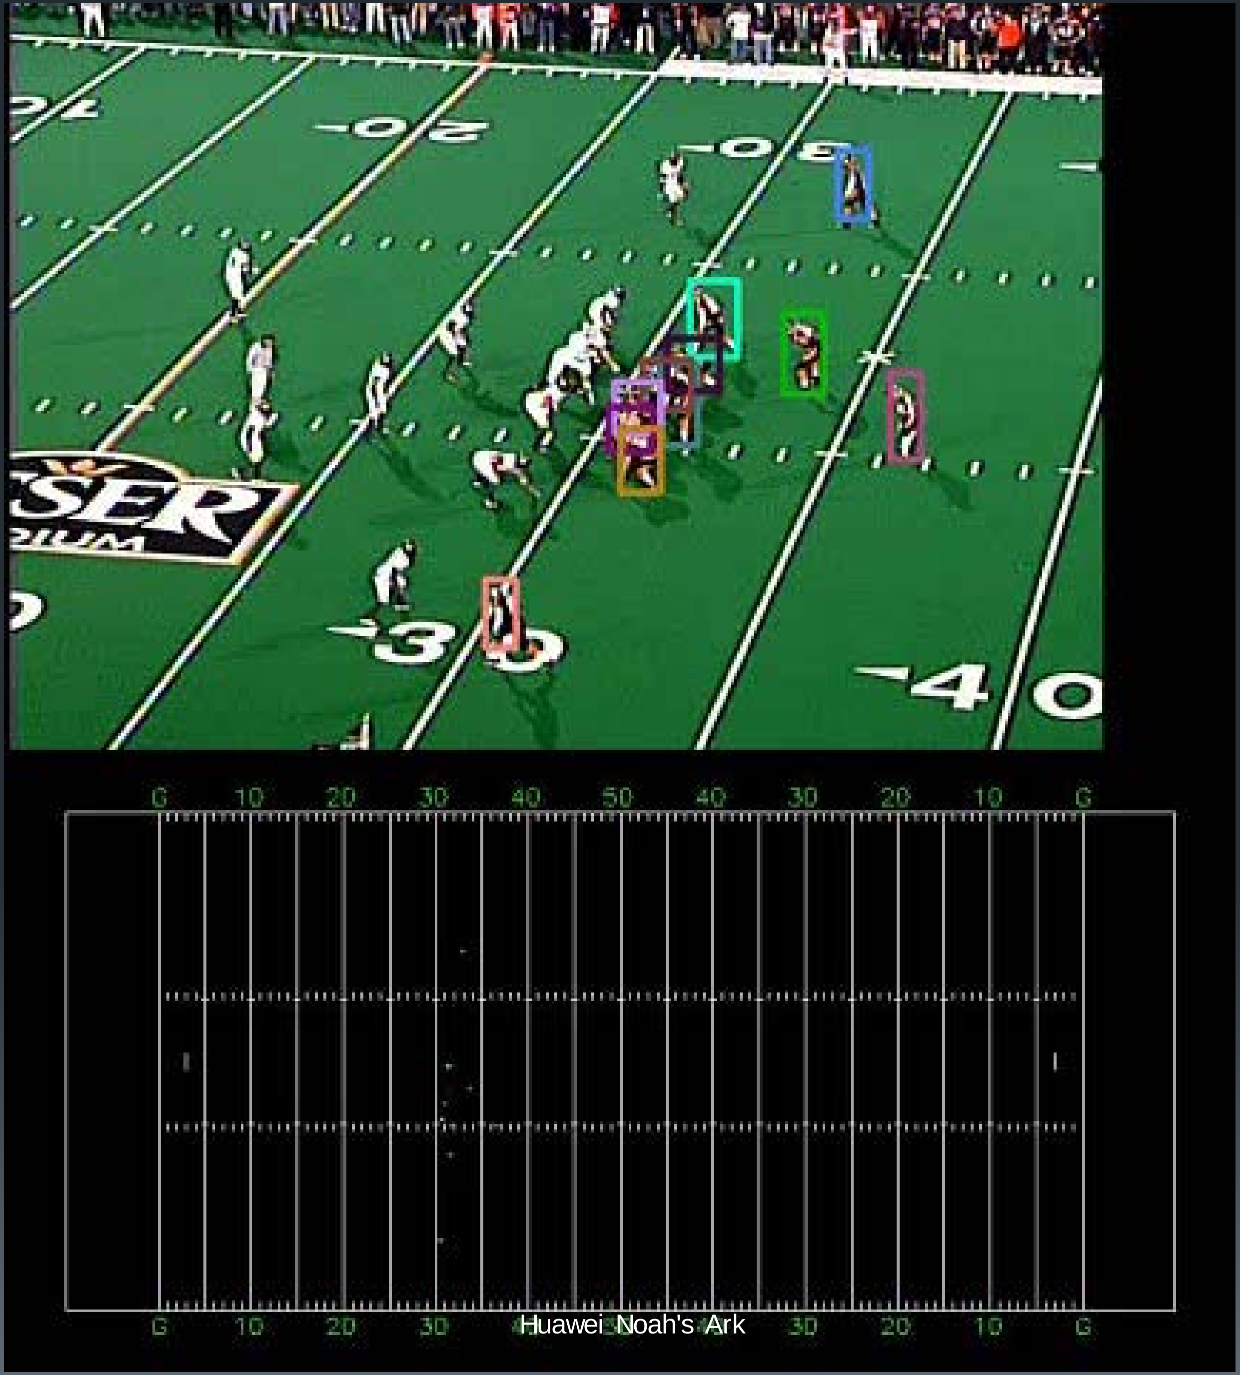
G 10 20 30 40 50 40 30 20 10 G

Birdsong Recognition

Xiaoli Fern
Raviv Raich

- At each site, a microphone records birds singing
- Goal: Determine what species are present at each site each day
- Problem: Labeling data is very time consuming
- Solution: Just label 10s segments
- Superset Label Problem

Segmented Spectrogram

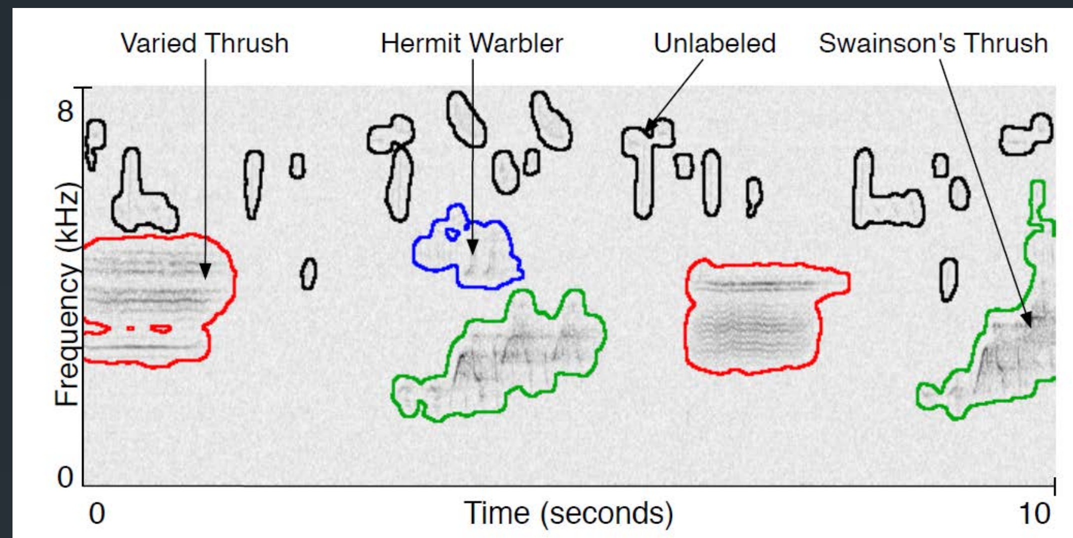

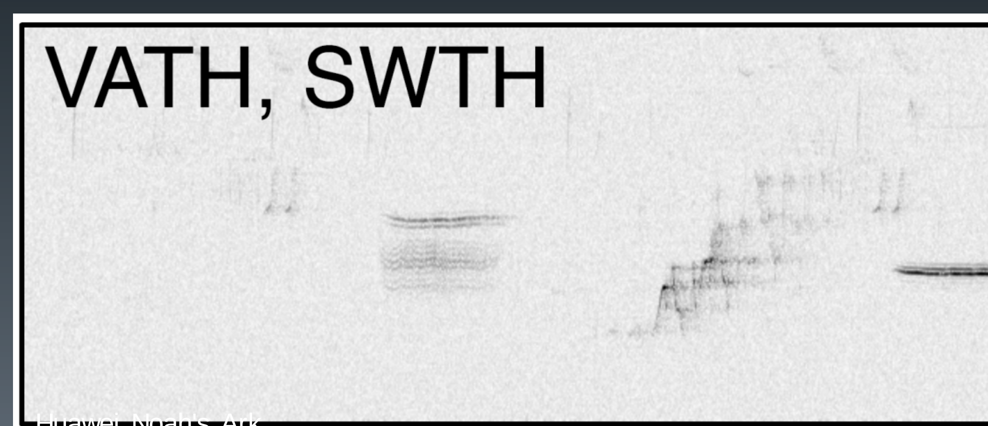

Machine Learning Approaches

- Method 1:
  - Formulate as "Multiple Instance Multiple Label" problem (MIML)
  - Each example is a "bag" of instances and is assigned multiple labels
- Method 2:
  - Formulate as "Superset Label Problem"
  - Each example is a single instance but has been assigned multiple labels

Results on Superset Labeling/MIML Problems

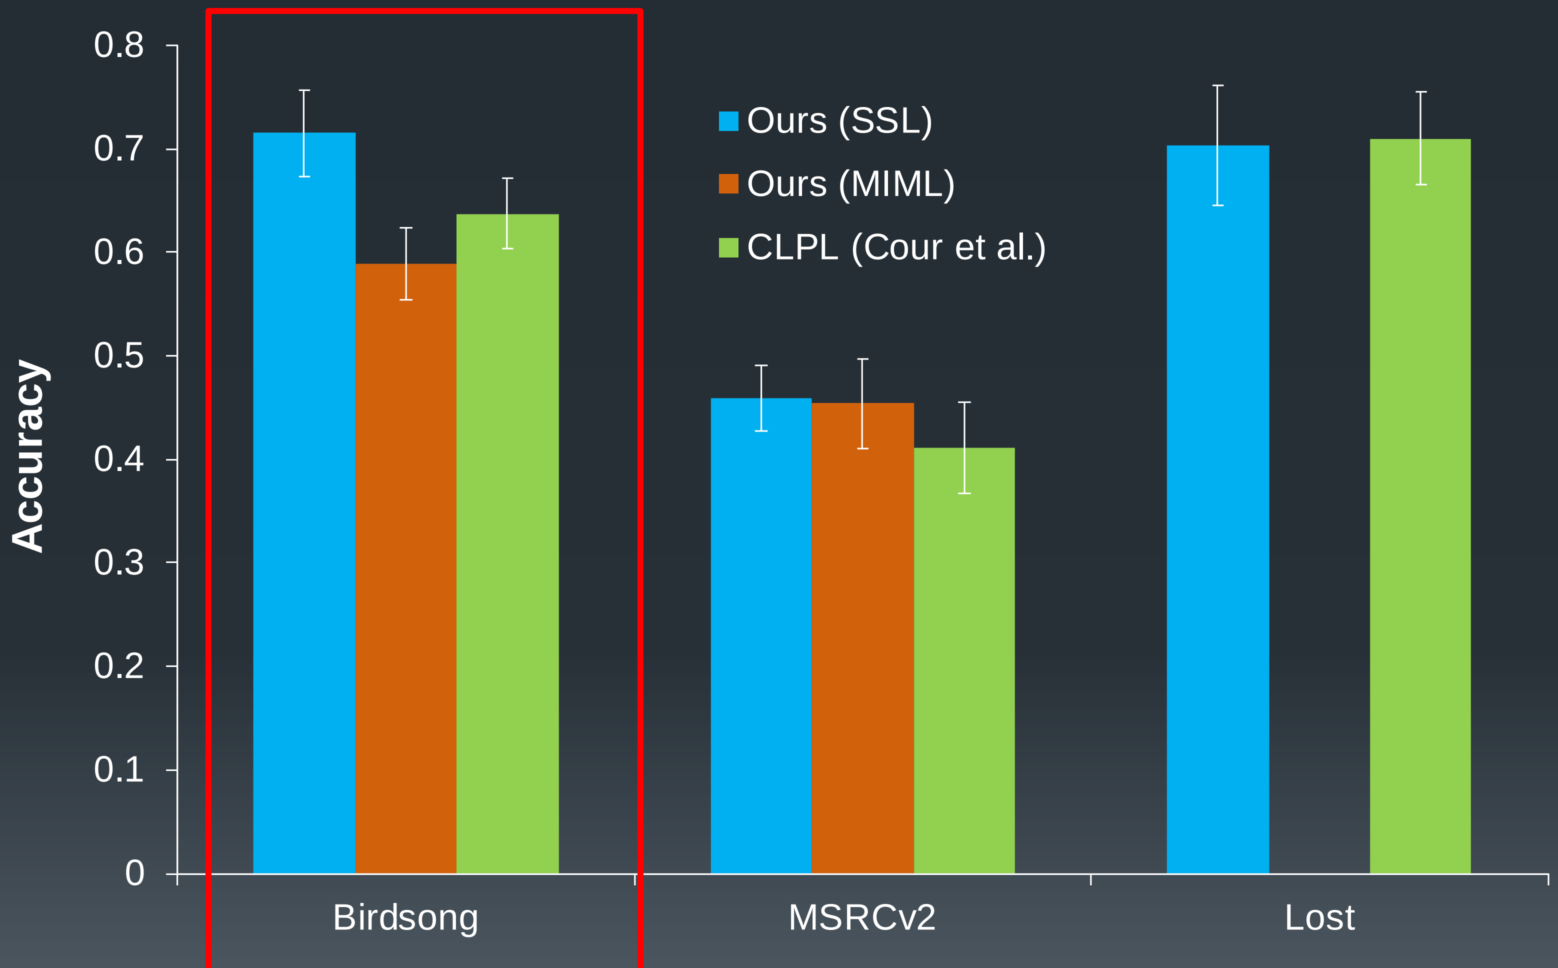

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Natural Language Processing

- Parsing
- Event and Event Argument Coreference

Liang Huang

Prasad Tadepalli

Xiaoli Fern

Linear Time Dynamic Programming Parsing

Liang Huang

- Very fast linear time dynamic programming parser
- explores exponentially many trees and outputs a forest
- state-of-the-art parsing accuracy in English and Chinese

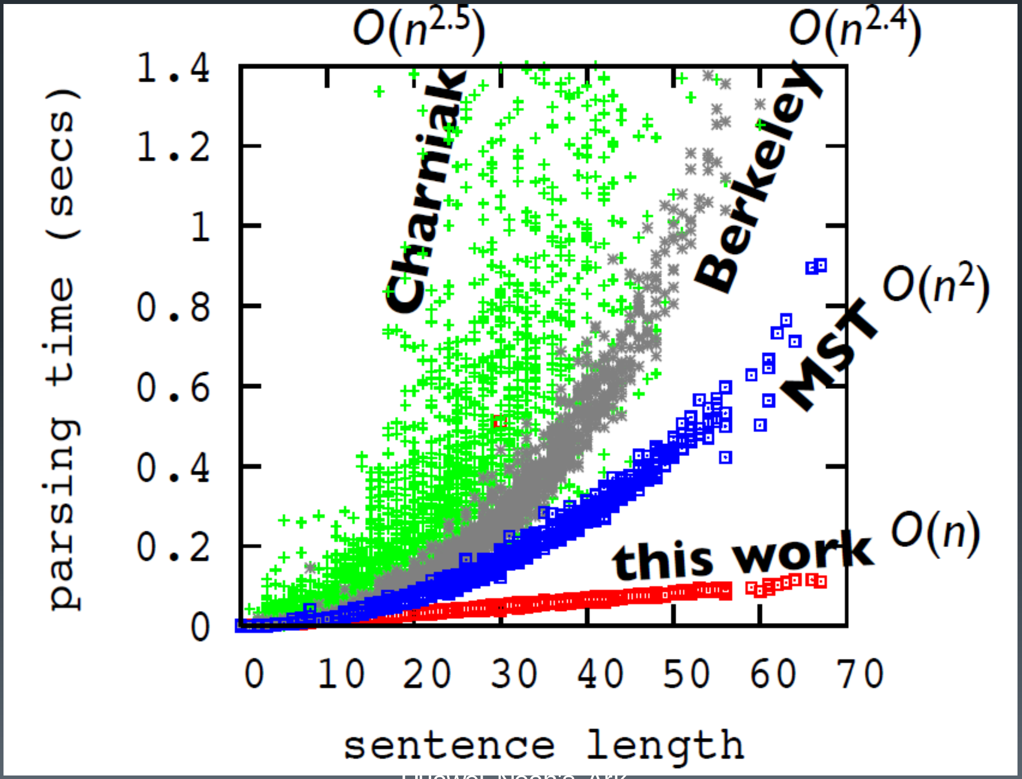

Joint Entity and Event Coreference

Xiaoli Fern
Prasad Tadepalli

- Two sentences:
  - "**Hugh Jackman plays a furry comic-book hero.**"
  - **"The Australian actor is playing a superhero.**"

- Coreferences:

Document 1	**Document 2**
"Hugh Jackman"	"The Australian actor"
"plays"	"is playing"
"a furry comic-book hero"	"a superhero"

Joint Inference using Learned Scoring Functions

- Learning performed using the "Easy First" framework
- Small but consistent improvement by using better learning method

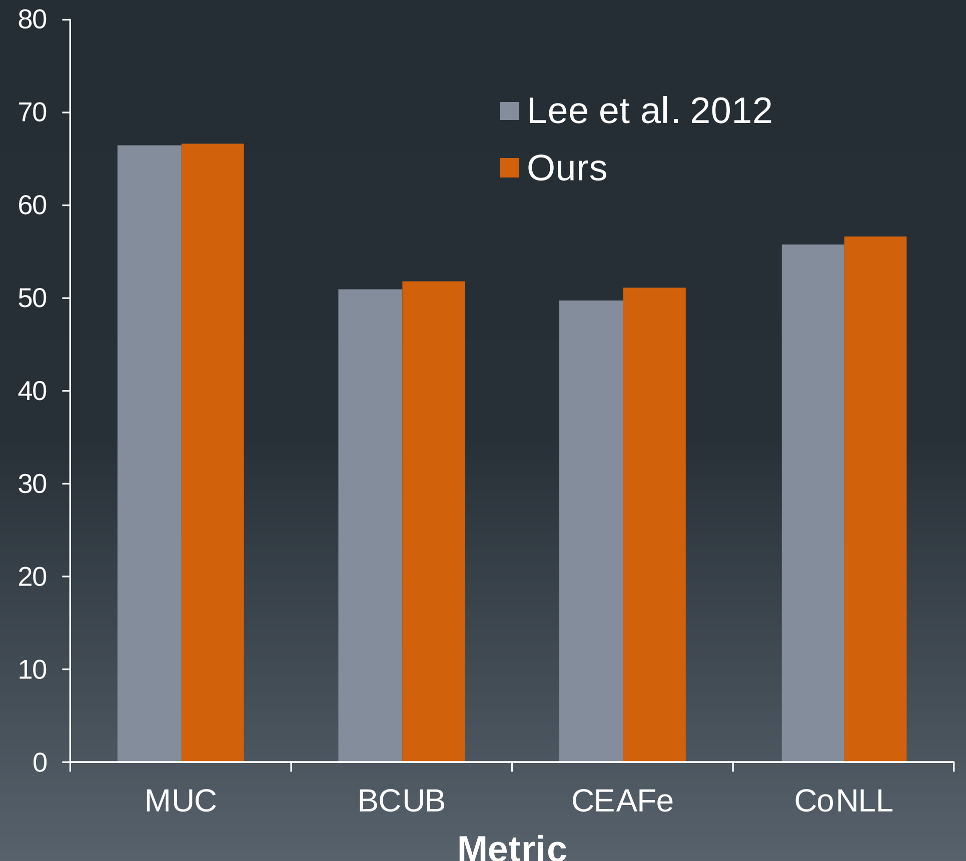

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Robotics

- Robotics PhD Program ranked #4 in US
- Home of ROS: Robot Operating System

Robot Locomotion: Legs

Jonathan Hurst

More Locomotion

- Ross Hatton
  - Snake robots and full-body locomotion
  - Casting manipulation (free cables and whip-like objects)
  - Spiders and spider webs
  - Vibration
- Belinda Batten
  - Micro-air vehicles
- Geoff Hollinger
  - Underwater Vehicles
  - Wheeled Vehicles

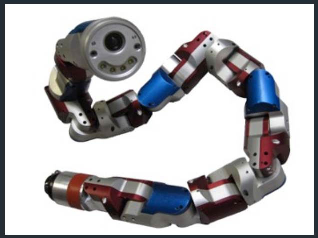




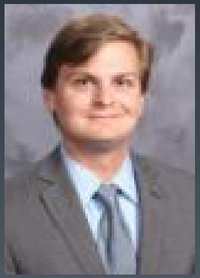

Oregon State University

Prosthetics and Human Control

- Robotic Control and Dynamics
- Human neuro-biomechanics
- Biomedical Implants
- Robot Hands

Ravi Balasubramanian

John Mathews

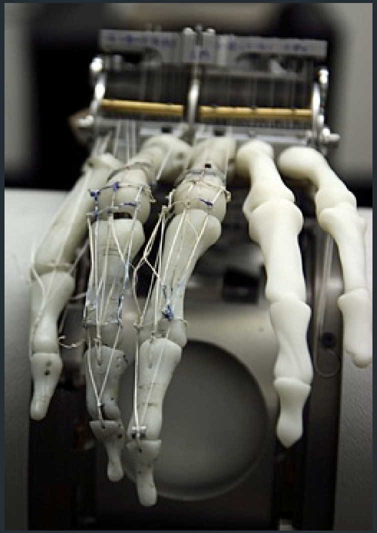

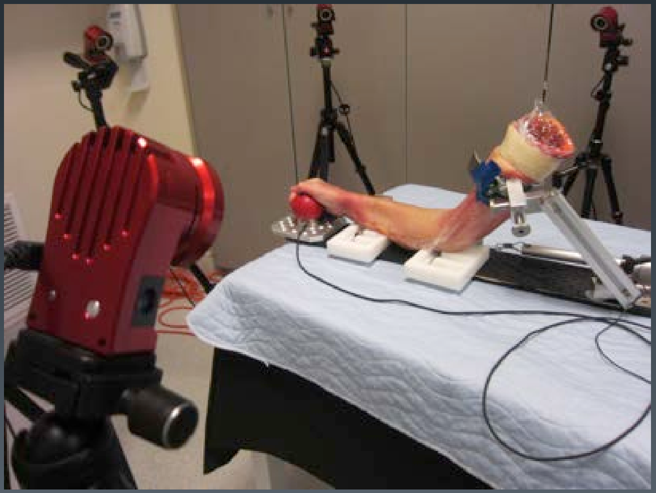

Oregon State University

Biologically-Inspired and Soft Robotics

- Soft sensors and actuators
- Gecko-inspired adhesion
- Wearable soft robotics

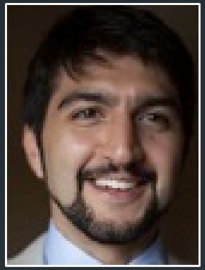

Yiğit Mengüç

Robot Teams and Autonomy

- Coordination of autonomous cars
- Air Traffic flow optimization
- UAV traffic management

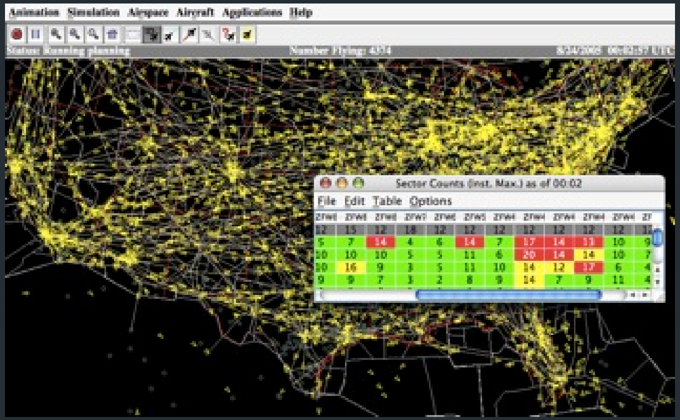

Kagan Tumer

Personal Robots, Autonomy, ROS

- Human-robot interaction
- Long-term robot autonomy
- Shared-autonomy human-robot systems
- Open-source software architectures for robotics
- Robots in the theatre
- Machine learning for the control of physical systems
- Advanced user interfaces for robot assistants

ROS

Bill Smart

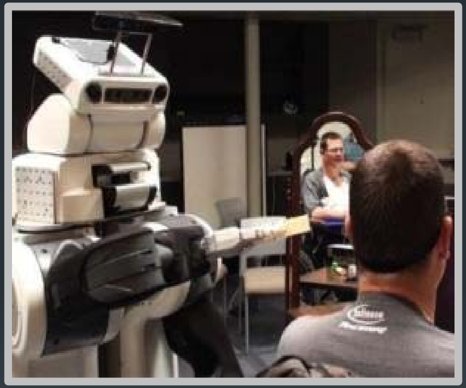

Oregon State University

CARIS
Center for Autonomy, Robotics, and Intelligent Systems

- Kagan Tumer, Director
- Alan Fern, Associate Director

- Artificial Intelligence
- Robotics
- Autonomy
- Human-AI and Human-Robot Interaction and Collaboration

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy

Questions?

Outline

- Machine Learning and Data Mining
  - Anomaly Detection
  - Data Cleaning
  - Robust ML
  - Bird Migration Modeling
- Planning and Control
  - Bayesian Optimization
  - Power Grid Control
  - Conservation Planning
- Computer Vision and Acoustics
  - Activity Recognition
  - Semantic Segmentation
  - Understanding Sports
  - Birdsong Classification
- Natural Language Processing
  - Parsing
  - Entity and Event Coreference
- Robotics
  - Locomotion: Legs
  - Snakes
  - Micro Air Vehicles
  - Underwater Vehicles
  - Hand Prosthetics
  - Wearables
  - Robot Teams
  - Personal Robots
  - Autonomy