AI 1-minute Madness

On 9/29/2014, the WSU EECS AI Group held the first annual AI 1-minute madness where all AI faculty and students gave a 1-minute overview of their research. Their slides are included here by group: Diane Cook, Aaron Crandall, Jana Doppa, Larry Holder and Matt Taylor. Thanks to Matt for organizing.

Activity Learning in the Wild

- Methods
  - Activity recognition
  - Activity discovery
  - Activity forecasting
- Applications
  - Understand behavior
  - Correlate with parameters of interest
  - Design activity-aware services

WASHINGTON STATE UNIVERSITY

CASAS

Center for Advanced Studies in Adaptive Systems

Diane J. Cook

cook@eecs.wsu.edu

Wearable Sensors for Ecological Rehabilitation

Gina Sprint, Vladimir Borisov, Diane Cook, Doug Weeks

- Investigating changes in movement over one week
- Inpatient rehabilitation participants
- Inertial measurement units on the waist and lower legs
- Comparison to reference population
- Prediction of clinical assessment scores

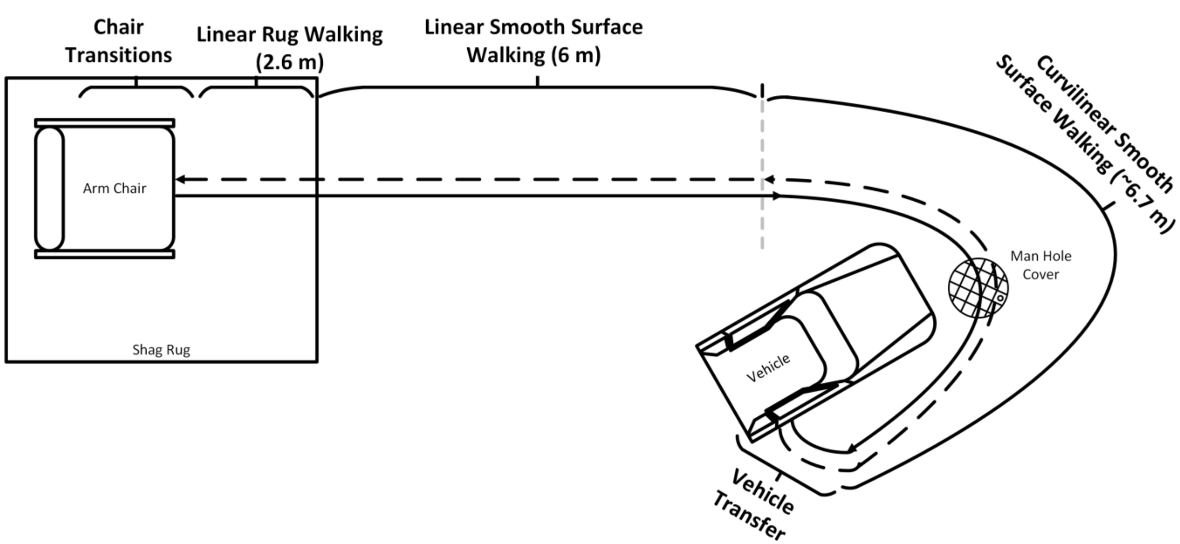

Figure 1. Participants perform the above ambulation circuit at two sessions, one week apart.

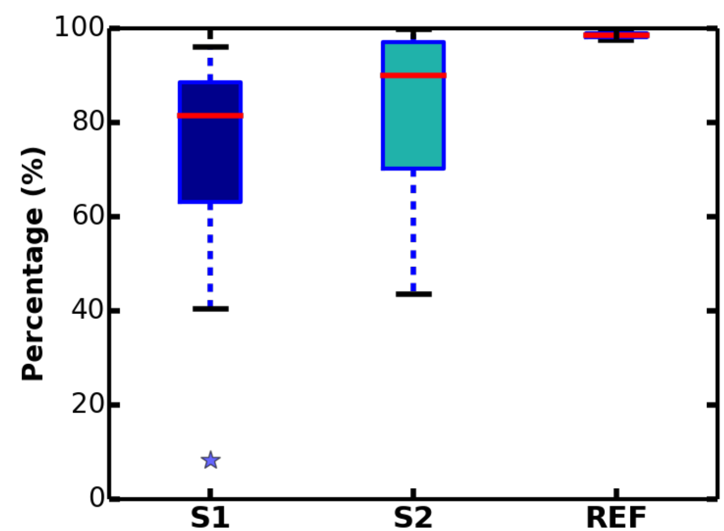

Figure 2. Step symmetry metric results.

Jessamyn Dahmen

"Digital Cognitive Assessments"

School of EECS, Washington State University

Advisor: Dr. Diane Cook

jessamyn.dahmen@email.wsu.edu

Research Areas: Gaming, Digital Interfaces, Cognition, Gerontechnology, Smart Environments

Cool Feature: A **nonintrusive** and more **objective** way to measure thinking, reasoning, and memory processes in **natural environments** over long periods of time

Paper and Pencil Psychological Tests

1

Tablet Games

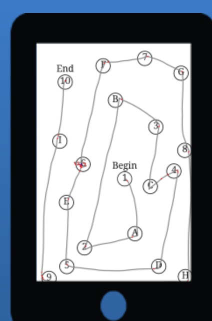

2

Digital Interfaces and Sensors Give Deeper Insights and More Data

3

Results: Am I Showing Signs Of Memory Loss?

Image References
1. www.iconarchive.com/
2. www.simpleicon.com/
3. www.betsmartmedia.com/

Automated Health Assessment using Smart Home Sensors

Prafulla N. Dawadi , Diane J. Cook

- Activities of Daily Living (ADL) = { Sleeping, Cooking, Eating, … }
- ***Declines*** in the ability to independently complete ADL are associated with shorter time to conversion to dementia and poorer quality of life.
- ***Monitor*** and ***detect*** changes in everyday ADL behavior using smart home sensors and learning algorithms.
- Use ADL information to ***perform*** health assessment.
- ***Keep*** older adults functioning in their own home for longer period than is currently possible.

Smart home

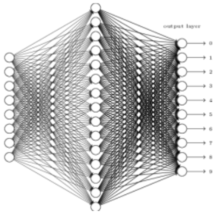

Learning algorithms

Monitor

Activities of daily living

Predict

HEALTH ASSESSMENT

Home Automation for Energy Efficiency

- Brian Thomas
- Context Aware
- Simulation
    - HMM

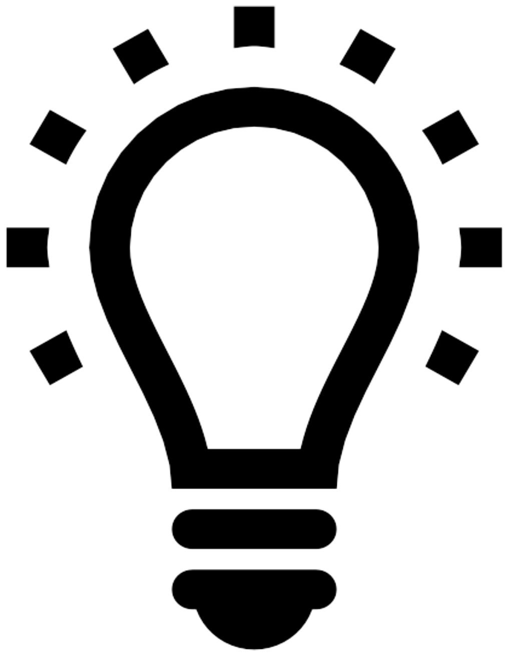

Christa Simon, B.S. (Psychology)

M.S. Psychology – anticipated May, 2015

Objectives

- To offer exceptional skills to identify and guide products from conception to launch.
- Professional, reliable, and well versed to understand and analyze user needs.

Skills

- Strong quantitative research background and experience in research design and analyzing datasets (SPSS, Mplus, Stata).
- Excellent organizational and analytical skills.
- Excellent communication skills, specifically in cross-functional settings.
- Interest in creating and analyzing products.
- Efficient in Microsoft Office (Excel, Word, PowerPoint).
- Working skill- technical abilities (Python, R)

Current Research Projects

- Smart Home: Activities of Daily Living, Executive Functioning and Aging / Role: Research Assistant
- Brain Health Intervention: Using wearable devices / Role: Project Lead
- Motivational adaptive rewarding using wearable devices / Role: Project Lead
- Digital Memory Notebook / Role: Project Lead

Future Goals

- To work for a tech startup company and gain knowledge in product development.
- Work with engineers in a cross-functional team to define products.
- Build and launch new products and features.
- Define product vision and future direction.

Salikh Bagaveyev

"Active Learning for Activity Recognition"
School of EECS
sal.bagaveyev@wsu.edu

Research Areas:
Smart Homes, Active Learning (Machine Learning), Crowd Sourcing

Cool Feature:
Reducing the burden of labeling data required for machine learning algorithms. Utilizing the annotators in a more efficient manner.

Smart Environments: Scalability, User Experience, Sensors Exploration, and Undergrads of Gerontech

Aaron S. Crandall, PhD.

Assistant Research Professor

Washington State University

WSU CASAS

http://casas.wsu.edu

Smart Environments and Multiple Residents

- Initial work on handling multiple residents:
  - Where are people?
  - Who is causing events?
- Use of statistical models to emit likely behaviometrics
- Sought to operate without wearables nor biometric solutions

Smart Environments Scalability

- Smart Home in a Box (SHiB)
  - Reducing installation complexity
  - Tested with pilot studies
  - Shipping to new participants
  - Establishing one of the world's largest public data gathering projects of this kind (80+ homes)
- Questions such as:
  - How do people perceive these systems?
  - How many sensors do we need per site?
  - How to install complex equipment?

Decision Support for Caregivers

- Smart Environments for eldercare and caregivers
- Evaluating what information caregivers need, when, and why?
- Caregivers:
  - √Residents
  - √Families
  - √Nurses
  - – Physicians
  - – Administrators

We are here

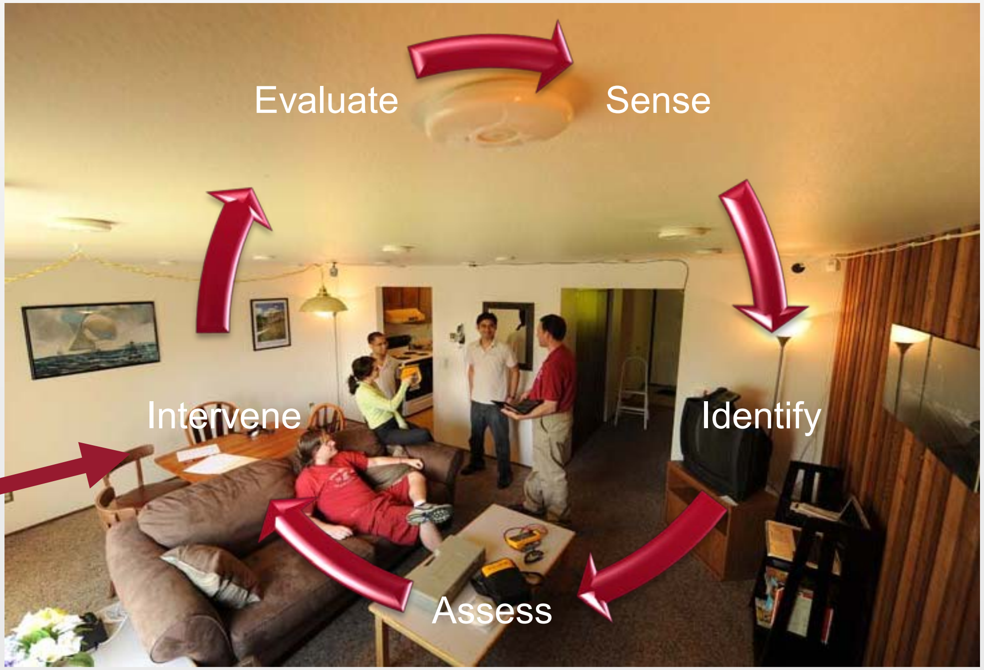

New Sensors Evaluation

- Current sensors:
  - PIR Motion, Lumen levels, temperature, contact, accelerometers, smart phones, power metering
- New sensors being tested and evaluated:
  - Kinect (v1 & v2)
  - Panasonic Grid Eye
  - Possible smart phone/mobiles (again)
  - Open to ideas!


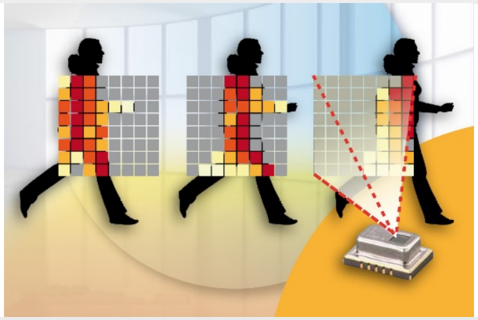


Undergrads in Gerontechnology

- Bringing undergrads into the field of gerontechnology (technology for aging)
- NIH Grant for:
  - Summer research programs
  - Senior capstone projects
    - Current one involves Kinect v2 and Parkinson's
  - Undergradute course series on gerontech
  - Direct engineering problems for CASAS




Learning and Reasoning for Data-driven Decision-making

Janardhan Rao (Jana) Doppa

School of EECS, Washington State University

Direction 1: Structured Prediction

- Large amounts of data
  - Text
  - Images
  - Videos
  - Scientific data

Low-level unstructured data

High-level Structured Representation

Structured Prediction: Examples

- **Sequence labeling**
  - POS tagging

    x = "The cat ran" y = *<article> <noun> <verb>*

  - Text-to-speech mapping

    x = "photograph" y = /f-Ot@graf-/

  - Handwriting recognition

    x =  y = s t r u c t u r e d

Y = set of all possible output sequences **Exponential !!**

Structured Prediction: Examples

- **Parsing**

x

"Red figures on the screen indicated falling stocks"

y

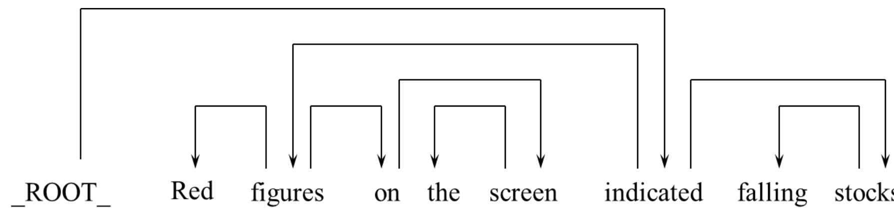

$\boldsymbol{Y}$ = set of all valid parse trees **Exponential !!**

- **Coreference Resolution**

x

*"**Barack Obama** nominated **Hillary Clinton** as his **secretary of state** on Monday. **He** chose **her** because **she** had foreign affair experience as a former **First Lady**."*

y

*"**Barack Obama** nominated **Hillary Clinton** as his **secretary of state** on Monday. **He** chose **her** because **she** had foreign affair experience as a former **First Lady**."*

$\boldsymbol{Y}$ = set of all possible clusterings **Exponential !!**

Structured Prediction: Examples

- **Scene Labeling**

x

y

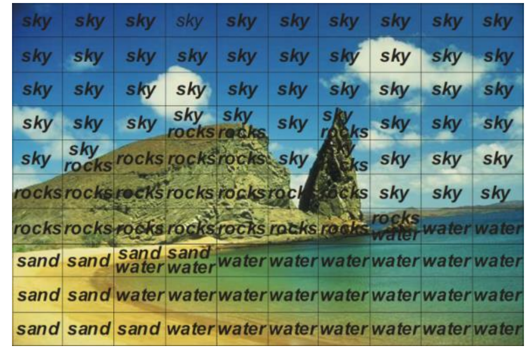

- **Biological Image Analysis**

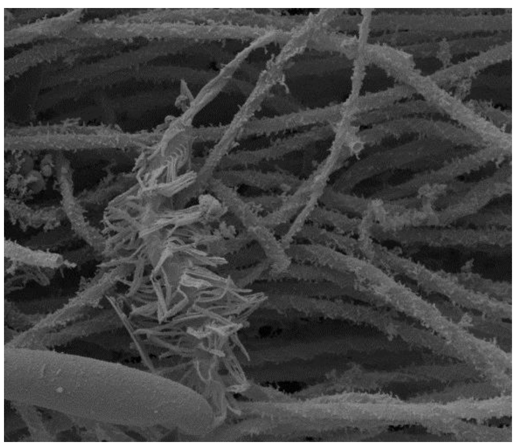

Nematocyst Image

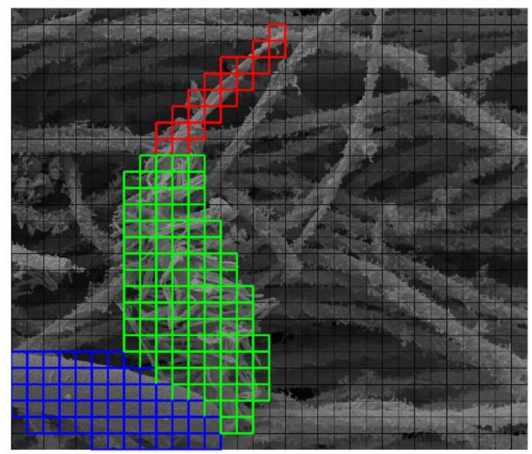

Body parts of the nematocyst

Y = set of all possible labelings **Exponential !!**

Direction 1: Structured Prediction

- **Structured Prediction**
  - Combinatorial optimization with unknown cost function
- **Learning to search the space of structured outputs**
  - How can we tradeoff speed and accuracy of making predictions?
  - Time-bounded inference algorithms?
  - Anytime inference algorithms?
- **Applications**
  - NLP, computer vision, bio-informatics, smart-home environments, and health informatics

Direction 2: Learning for Speedup and Resource-bounded Optimization

- **Learning to Speedup combinatorial optimization**
  - Given a cost function, how can we learn control knowledge (e.g., heuristics, pruning rules) to find solutions quickly without loss in accuracy?
  - Problem dependent optimization?
- **Resource-bounded optimization**
  - Optimization under resource-constraints (e.g., time, memory, information …)?
- **More generally**
  - Bounded rationality? Provably bounded optimality?

Direction 3: Data-Driven Decision-Making

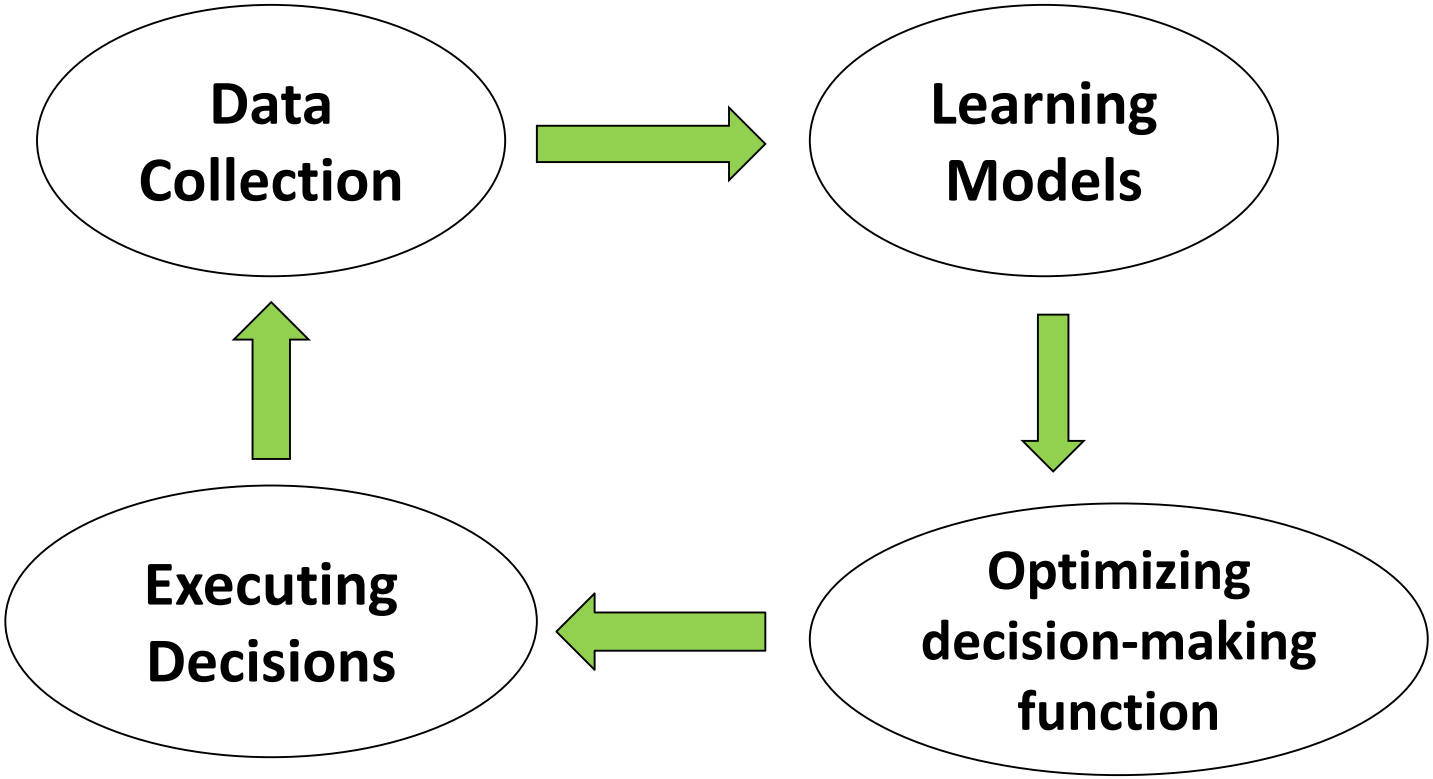

- **Closing-the-loop**
- No "one-size-fits-all" solution
- Identify a set of similar problems, and find solutions

Direction 4: Learning from Weak Supervision

- **Generic and principled approaches for**
  - Incorporating prior-knowledge into learning algorithms
  - Semi-supervised and Multi-view learning
  - Active learning
  - Learning from noisy supervision (e.g., crowd-sourced data)
  - Learning from rich-forms of supervision
  - Human-in-the-loop learning

Questions ?

Smart Environments Research Center (SERC)

Larry Holder, Director

Problem

- Making the environments in which we live and work safer, healthier and more productive.

Hypothesis

- Advanced data analytics
- Adaptive systems
- WSU a world leader in smart environments research

Approach

- Multi-disciplinary research
  - AI, Machine Learning, Robotics, Data Science, Pervasive Computing
  - Psychology, Neuroscience
- Education
- Outreach
- Technology development and transfer

Results

- SERC approved by WSU faculty senate (Feb 2013)
- SERC website (Sep 2014)
- SERC NSF Research Experience for Undergraduates (REU) under review
- NSF Computing Research Infrastructure Proposal (Smart Campus)

http://serc.wsu.edu

Pattern Learning and Anomaly Detection in Streams (PLADS)

Larry Holder, Professor

Problem

- Find patterns and anomalies in streaming data
- Fuse data from multiple, heterogeneous sources
- Scale to large data sizes and stream rates

Hypothesis

- Represent data as dynamic graph
- Graph mining for patterns and anomalies
- Partitioning and sampling of stream(s) is sufficient

Approach

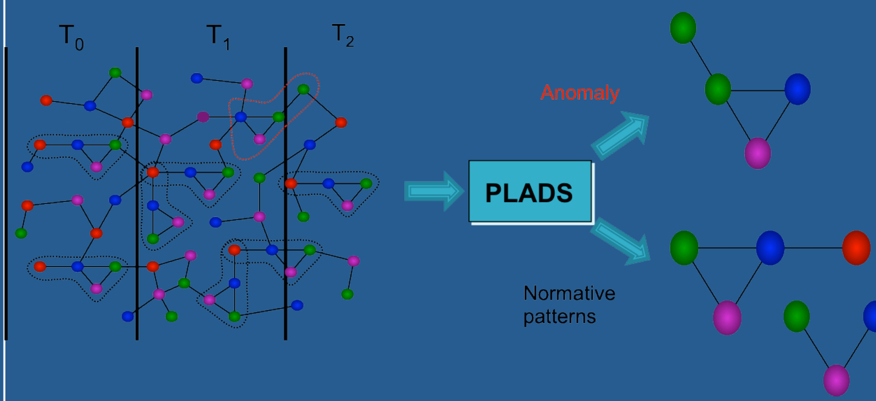

Results (e.g., Twitter)

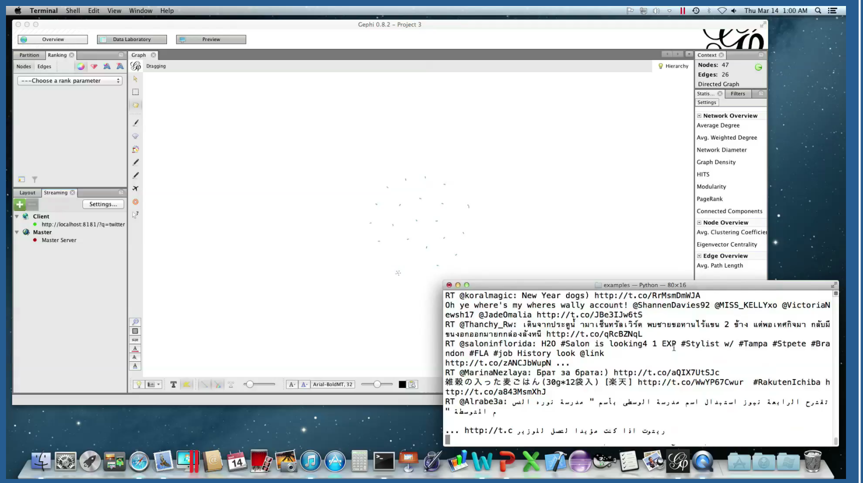

www.ailab.wsu.edu/adgs

Applying Graph Mining to Wireless Sensor Network Data

Name: Syeda Selina Akter Advisor: Larry Holder

Problem

- Improve performance on recognition and prediction tasks for various wireless sensor networks

Hypothesis

- Graph representation and mining can achieve performance improvement

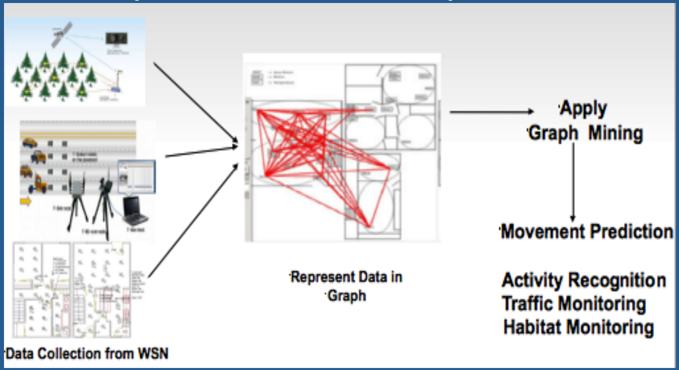

Approach

- Smart home and smart phone sensors
- Features: Count of edges/subgraphs

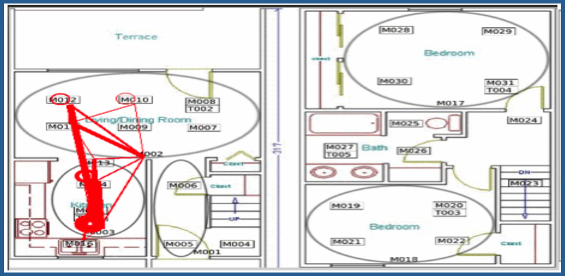

Graphical features for activity “watch TV” from Tulum testbed

Results

- Improvement in activity recognition accuracy compared to baseline in three different environments

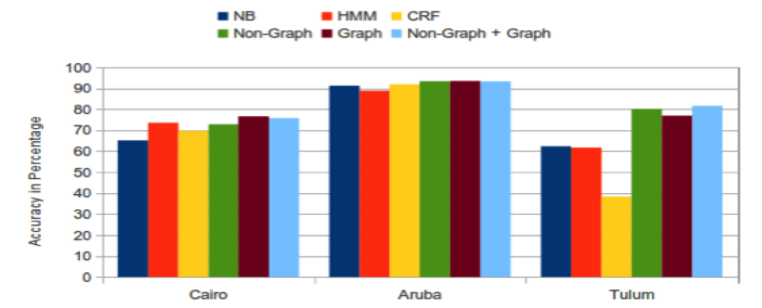

Reference : Akter, Syeda, and Holder, Lawrence B. “Activity Recognition using Graphical Features”. Machine Learning and Applications, 2014. ICMLA '14. 13th International Conference on. IEEE, 2014.

Lossy Compression of Graph Data

Name: Jason Fairey Advisor: Larry Holder

Problem

- Identifying instances of patterns in large or streaming graphs

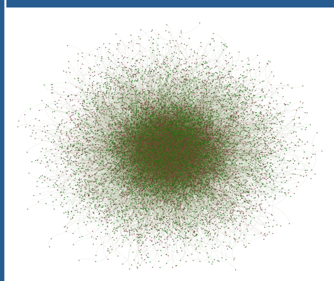

- Data comes in faster than can be processed using strict methods.

Hypothesis

- Lossy compression of graph information can identify instances of given patterns.
- Lossy methods allow for efficiency with manageable cost to accuracy.

Approach

- Use compression method similar to the way the human retina works.
- Layers of compressions.

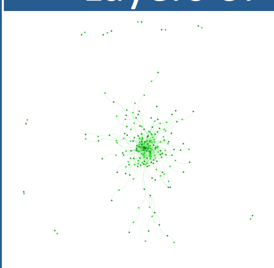

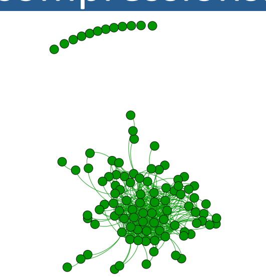

Layer	Nodes	Edges
0	11000	18414
1	211	532
2	100	528

Results

Nodes	Edges	Pattern	Time (Static)	Time (Stream)	Accuracy
11000	18414	10/3	23ms	180ms	100/100
102500	156997	25/4	290ms	2753ms	103/100
1025000	1582498	250/5	11475ms	69845ms	111/100
1001000	1521376	10/3	989ms	124218ms	101/100

Single Trials
Single core in debug environment

Frequent Subgraph Discovery in Large Attributed Streaming Graphs

Name: Abhik Ray Advisor: Larry Holder

Problem

- Find frequent subgraphs in streaming graph data.
- Scale to large data size and stream rate

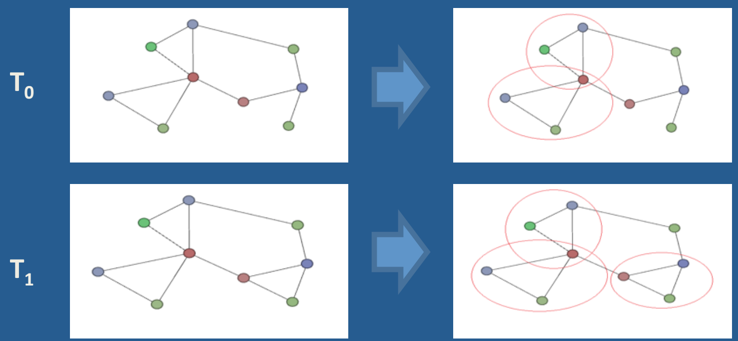

Hypothesis

- Changes to the set of frequent subgraphs will come from the neighborhood of the newly added nodes and edges.

Approach

- Sample graph neighborhood around the newly added nodes/ edges from stream.
- Convert to graph transactions and frequent subgraph mining for graph transactions.

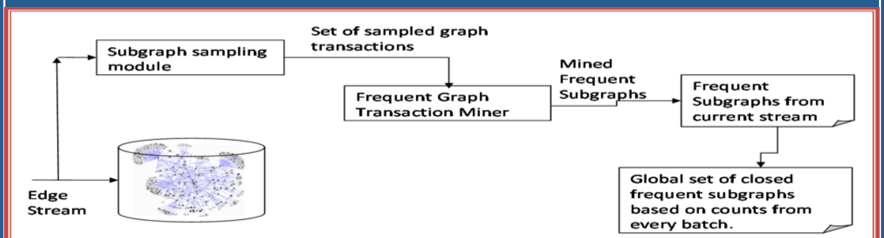

System overview of Frequent Subgraph Miner for Streaming Large Graphs
(large network picture taken from http://www.facetracks.com/wp-content/uploads/2011/08/SocialNetworkAnalysis_Graph.jpg)

Results

- 100% accuracy on known patterns from artificial graph.
- Faster than stream-rate.

Table 1: Maximum running time vs. Evolution Time

Dataset	Evolution Time	Maximum Running Time
Twitter	1 year	4,817 seconds
Hetrec	98 years	12, 939 seconds

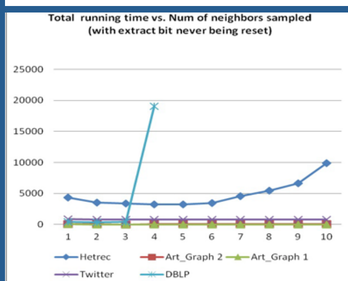

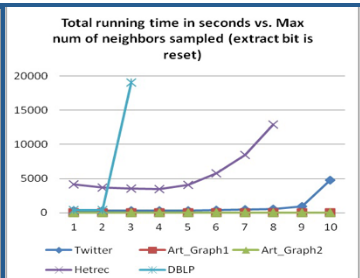

*A. Ray, L.B. Holder, S. Choudhury, "**Frequent Subgraph Discovery in Large Attributed Streaming Graphs**". KDD 2014, New York: 3rd International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications.*

Supervised Learning in Dynamic Networks

Name: Yibo Yao Advisor: Larry Holder

Problem

- Find discriminative patterns in large and dynamic graphs
- Update classification model quickly when new data stream in

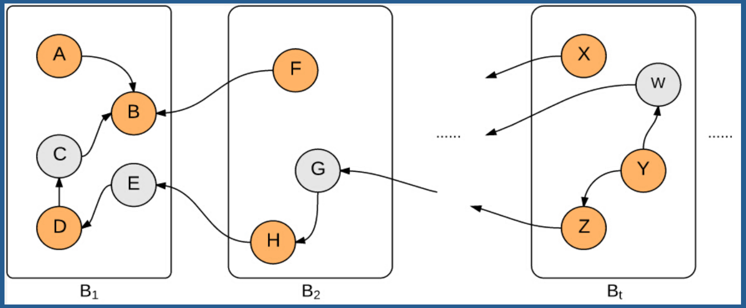

Approach

- Extract subgraphs
- Classify subgraphs using graph kernels plus SVM
- Incrementally update a learner using support vectors

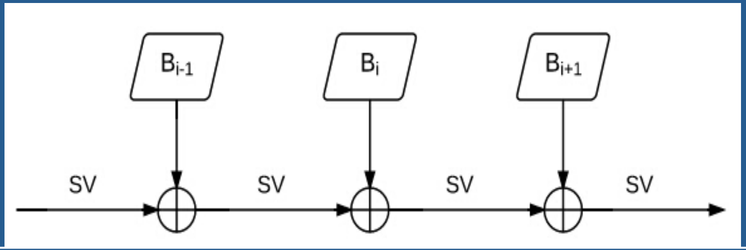

Hypothesis

- Graph representation:
  - node → entity
  - edge → relation
- Interconnected entities have similar properties
- Explore interconnection using graph mining to facilitate classification

Results

- Reduced memory usage and training time
- Improved average classification accuracy by 5-10%

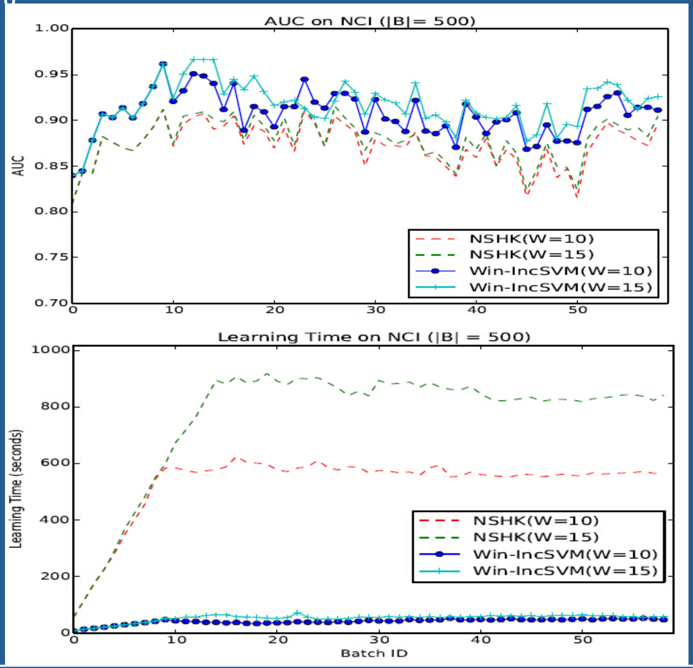

1. Y. Yao and L. Holder, Scalable SVM-based Classification in Dynamic Networks, 2014 IEEE International Conference on Data Mining (ICDM), to appear

Matthew E. Taylor

The Intelligent Robot Learning Laboratory

Washington State University

Robotics Club

RoboSub

Transfer for Reinforcement Learning

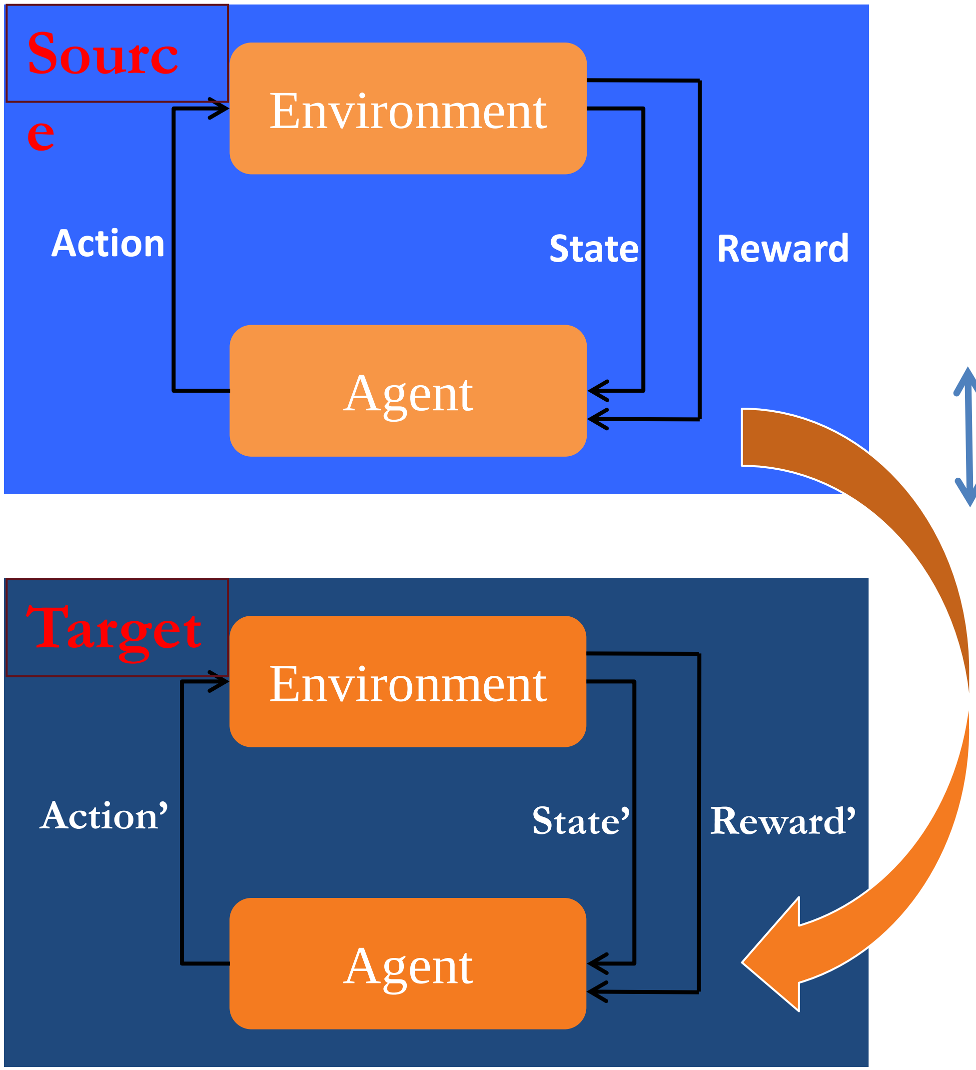

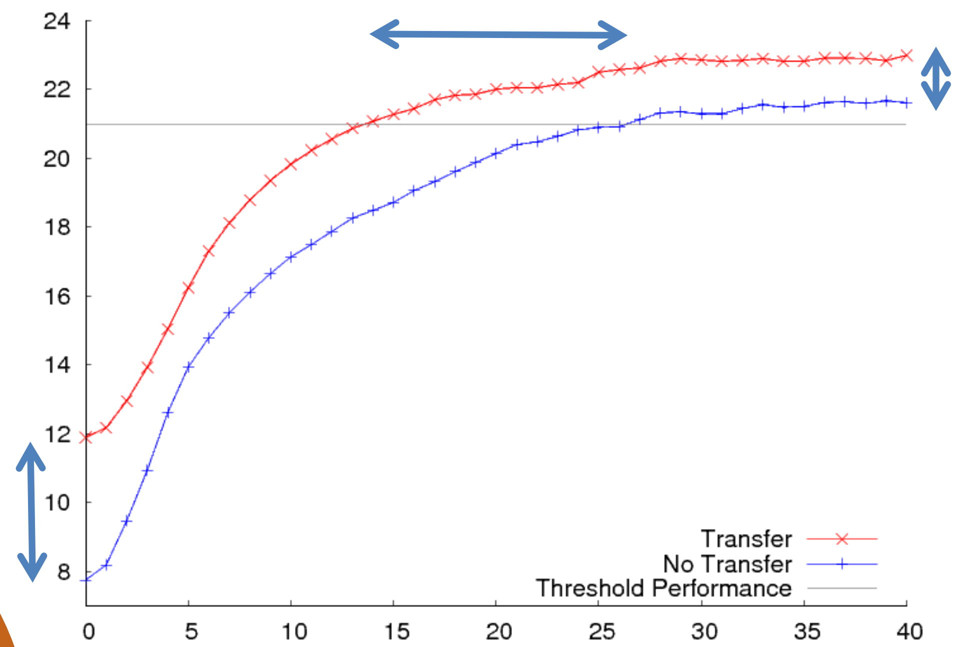

http://eecs.wsu.edu/~taylorm

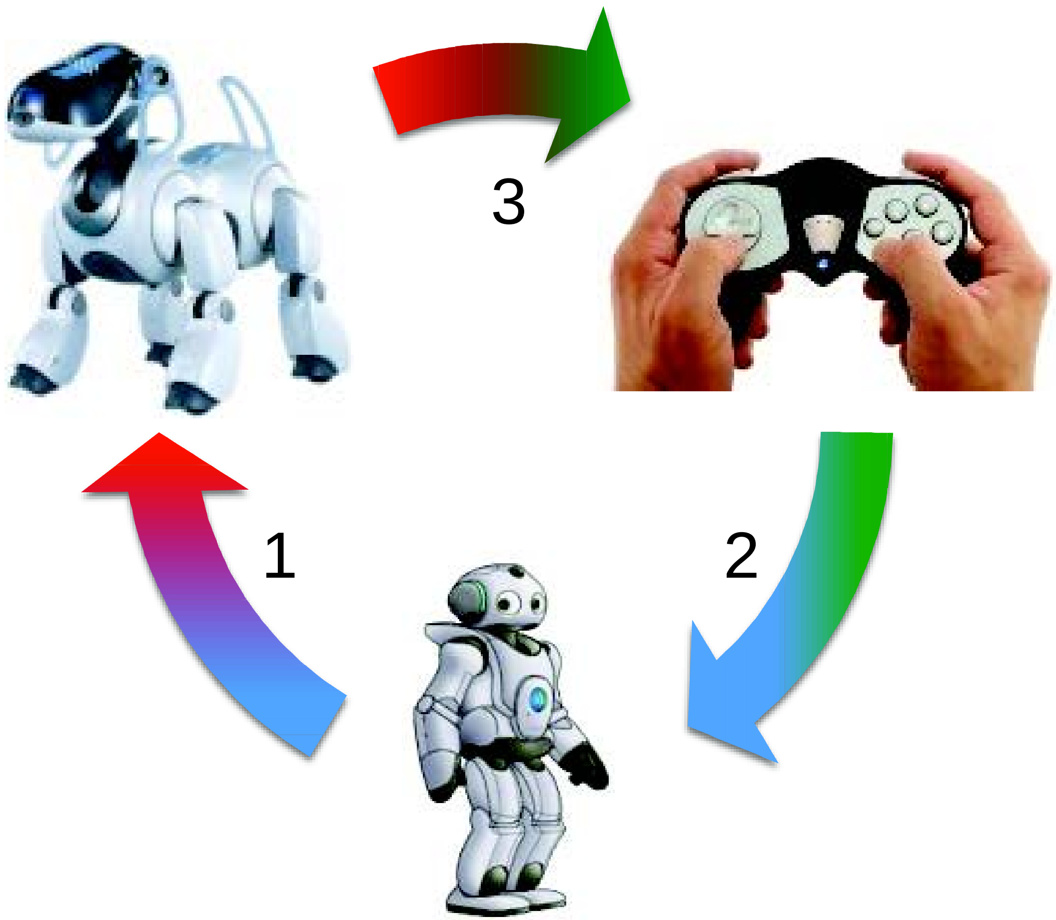
3
1
2

Yusen Zhan

- Project Name: Agent Teaching Humans
- Agent teaching humans to play video games
- Developed new teaching algorithms
- Future: More experiments

Yusen Zhan

- Project Name: On-line Transfer Learning for RL
- Completed work
    - Develop a formal framework for teaching/advising
    - Convergence Analysis
- Future: Machine/human teachers or human/machine students

A Strategy-Aware Technique for Learning Behaviors from Discrete Human Feedback

Bei Peng   Matthew E. Taylor

Reward-Focused
Punishment-Focused
Balanced Feedback
Inactive

user study

avatars of agent

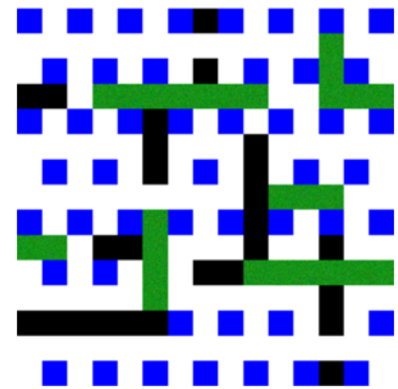

sequential domain

training conditions

A Strategy-Aware Technique for Learning Behaviors from Discrete Human Feedback

- Human trainers follow multiple strategies when training agents
- Explicitly considering trainer strategy can allow a learner to make inferences from cases where no feedback is given
- Training conditions affect trainers' choice of strategies
- Papers:

Robert Loftin, **Bei Peng**, James MacGlashan, Michael Littman, Matthew E. Taylor, David Roberts, and Jeff Huang. Learning Something from Nothing: Leveraging Implicit Human Feedback Strategies. *In Proceedings of the 23rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, August 2014.
Nominee for "RSJ/KROS Distinguished Interdisciplinary Research Award"

Robert Loftin, **Bei Peng**, James MacGlashan, Machiael L. Littman, Matthew E. Taylor, Jeff Huang, and David L. Roberts. A Strategy-Aware Technique for Learning Behaviors from Discrete Human Feedback. *In Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, July 2014. 28% acceptance rate

Agent Corrections to Pac-Man from the Crowd

Gabriel de la Cruz, Bei Peng, *Walter Lasecki and Matthew Taylor

- Using crowdsourcing to speed up agent learning, especially in the reinforcement learning paradigm.
- The long-term goal is to leverage the human crowd's advice to improve learning when and where it is needed.
- Currently, we are working on a paper with data from a user study that requires the crowd to suggest time of mistake

*Rochester University

Autonomous Unmanned Aerial Systems for Bird Deterrence

Problem

Solution

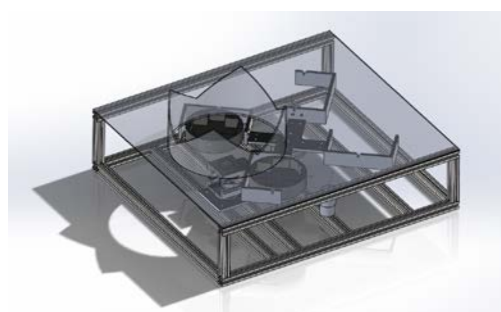

SI

Lifelong Learning for Heterogenous Robot Teams

Gabriel de la Cruz, Lorin Vandegrift, Dustin Crossman, Matthew Taylor,
University of Pennsylvania, and Olin College

Search & Rescue

Technologies used

- Parrot A.R. Drone
- Turtlebot
- Microsoft Kinect
- OpenCV
- Gazebo
- ROS (Robot Operating System)
- SLAM
- Reinforcement Learning
- PG-ELLA

Curriculum Development for Transfer Learning in Dynamic Multi-agent Settings

Anthony D. La , Matthew E. Taylor

- The Department of Computer Science at The University of Texas at Austin.
- ***High requirement*** of training experience to handle complex domains.
- ***Transfer learning*** helps but current methods involve manual definition.
- Create a novel framework for ***curriculum development for transfer learning***.
- ***Automate*** process of assigning easy (fast to learn) source tasks which still promote new, crucial behaviors.
- ***Evaluation*** in at least 2 domains.

StarCraft Target Task

Curriculum Generator

Strong:

- ✓ Shooting when under attack

Passing:

- ✓ Controlling choke point

Needs work:

- Hunting fleeing enemies
- Fending off enemies attacking the turret.

Specially Tailored Tasks

Learn Something New

Transfer Between Tasks

StarCraft Target Task

"Well done, android. The Enrichment Center once again reminds you that Android Hell is a real place where you will be sent at the first sign of defiance."