Conference on Automated Learning and Discovery

Workshop W1, Thursday, June 11th
Rangos 1, University Center

Learning Causal Bayesian Networks

Chaired by: Richard Scheines and Larry Wasserman

This workshop will focus on several topics in an area of automated learning and discovery that has flourished in the last several years: causal discovery.

Session I – Foundations

11:00-11:25:
"Graphical Representation of Causal Models"
Phil Dawid

11:25-11:45: Discussion

11:45-12:10:
"Limits on Learnability of Causal Hypotheses"
Jamie Robins, Richard Scheines, Peter Spirtes and Larry Wasserman

12:10-12:30: Discussion

12:30-2:00: Lunch

Session II – Learning Bayesian Networks

2:00-2:25:
"Learning the Structure of Dynamic Probabilistic Networks"
Nir Friedman, Kevin Murphy and Stuart Russell

2:25-2:45: Discussion

2:45-3:10:
"A Multivariate Discretization Method for Learning Causal Bayesian Networks"
Stefano Monti and Greg Cooper

3:10-3:30: Discussion

3:30-4:00: Coffee Break

Session III – Models for Causal Networks

4:00-4:25:
"Graphical models with interpretable structure"
Thomas Richardson

4:25-4:45: Discussion

4:45-5:10:
"Graphical Models and Exponential Families"
Chris Meek and Dan Geiger

5:10-5:30: Discussion

Conference on Automated Learning and Discovery

Workshop W2, Thursday, June 11th
Rangos 2, University Center

Mixed-Media Databases

Chaired by: Shumeet Baluja, Christos Faloutsos, Alex Hauptmann and Michael Witbrock

This workshop is intended for researchers with an interest in learning from multiple media. The workshop will emphasize both algorithms and applications of learning with mixed media databases. Papers to be presented in this workshop describe novel approaches for wide variety of applications for learning in mixed media databases, ranging from video indexing, tutoring systems, and medical case retrieval.

Session I

11:00 – 11:30: Invited Talk
"Presence: Accessing Live Multimedia Information Systems"
Ramesh Jain

11:30 – 12:30:
"Multiple Media Database System for TV Newscasts and Newspapers"
Watanabe, Okada, Kaneji, Sakamoto
"Information Retrieval by Image Content and its Associated Text in a Digital Library"
Meza, Favela
"Generating a Large Database of Named Faces in News Video" *Houghton*

12:30 – 2:00: Lunch

Session II

2:00 – 2:30: Invited Talk
"Automated Visual Search Tools Based on Integration of Multi-Level, Multi-Media Features"
Shih-Fu Chang

2:30 – 3:30:
"A Hidden Markov Model Framework for Video Segmentation Using Audio and Image Features"
Boreczky & Wilcox
"When Listening is Not Enough: Potential Uses for Vision for a Reading Tutor that Listens"
Kominek, Aist, Mostow
"Classification Based Medical Case Retrieval: Using Radiology Images and Patient Text Information"
Liu, Dellaert, Rothfus

3:30 – 4:00: Break

Session III

4:00 – 4:20:
"Discrete MMI Signatures for Media Classification and Clustering"
Foote

4:20 – 5:30:
Discussion for NSF Future Challenges and Opportunities for Mixed-Media Databases

Invited talks are 20 minutes, with 10 minutes for questions. Presentations are 15 minutes, with 5 minutes for questions

Conference on Automated Learning and Discovery

Workshop W3, Thursday, June 11th
Connan Room, University Center

Machine Learning and Reinforcement Learning for Manufacturing

Chaired by: Sridhar Mahadevan and Andrew W. Moore

In recent years there has been a flurry of research in statistics and machine learning on the problem of decision making and control. Exciting progress has been made in many areas, such as Reinforcement Learning, Neural Networks, and Bayesian Networks. State-of-the-art applications are emerging in the control of continuous processes, batch processes (such as wafer fabrication), probabilistic diagnosis, and industrial engineering tasks such as optimal control of transfer lines or production scheduling. The workshop talks have been selected to be of interest to a mix of researchers from industrial labs, university researchers who work on manufacturing problems, reinforcement learning researchers, and others in the fields of statistics and machine learning who want to learn more about potential machine learning applications in industry. The workshop will provide a forum for researchers in areas related to this topic to exchange ideas and debate recent case studies of machine learning and probabilistic modeling applied to manufacturing systems. We will also explore advances in algorithms or theory that permit learning approaches to scale up to bigger problems.

11:00 – 12:30:
Session I – Machine Learning for Manufacturing
"Parallel Q-learning for communication networks"
A. Nowe
"General Paradigm for Applying ML in Automated Manufacturing processes"
W. Zhang and R. Tjoelker
"Designing experiments for Learning with NNs"
O. Gür Ali et al.

12:30 – 2:00: Lunch

2:00 – 3:30:
Session II – Value Function Search Methods
"Value Function based production scheduling"
J. Schneider
"Learning Multi Buffer Sequencing"
S. Mahadevan et al
"RL for realistic manufacturing processes"
S. Hagen and B. Kröse
"Average Reward RL for Optimizing factories"
P. Tadepalli

3:30 – 4:00: Break

4:00 – 5:30:
Session III – Parametric Policy Search Methods
"Autonomous Response Surface Methods"
A. Moore et al
"High Quality Thermostat Control"
M. Riedmiller
"RL for intelligent building control"
J. Schulte

Each talk will be assigned 15 mins for presentation and 5 mins for questions. The remainder of the time in each 90 minute session will consist of round-table discussion.

Conference on Automated Learning and Discovery

Workshop W4, Thursday, June 11th
Dowd Room, University Center

Large-Scale Consumer Databases

Chaired by: Mike Meyer, Teddy Seidenfeld and Kannan Srinivasan

Large-scale consumer databases provide the opportunity for novel approaches to marketing and customer services.

Areas that we will address include:

- real-time customer profiling, for example, to detect fraudulent use of credit cards;
- longitudinal customer profiling to develop selective marketing strategies based on
- the life-cycles of customers;
- interactive surveys, with adaptive (sequential) components for increasing relevant
- services to particular customers, e.g., automated delivery of commodity bundles

11:00 – 12:30: Session I

"A Formal Statistical Approach to Collaborative Filtering"
Lyle Ungar, Dean Foster

12:30 – 2:00: Lunch

2:00 – 3:30: Session II

"Automated Discovery of Discriminant Rules for a Group of Objects in Databases"
Tae-wan Ryu, Christoph Eick

3:30 – 4:00: Break

4:00 – 5:30: Session III

"Mining Data to Target market Consumers From a Large Relational Database"
Peter Johnson, Jim Delaney

Conference on Automated Learning and Discovery

Workshop W5, Friday, June 12th
Connan Room, University Center

Visual Methods for Massive Data

Chaired by: Bill Eddy and Steve Eick

This workshop is intended for researchers interested in the development and use of visual methods for the exploration and analysis of very large data sets. Visual methods include both static and dynamic displays. Very large data sets can include both standard numerical data and other complex and multi-media data sets such as text and images.

Session I
11:00 – 11:45

"A Beautiful Diamond has Many Facets"
Graham Wills, Bell Labs

11:45 – 12:30: Discussion

12:30 – 2:00: Lunch

Session II
2:00 – 2:45

"The Interplay of Large-Scale Data Analysis and Data Visualization"
Audreas Buja, AT&T Labs

2:45 – 3:15: Position paper

"When Speech Input is Not an Afterthought: A reading Tutor that Listens"
Jack Mostow and Gregory Aist

3:15 – 3:30: Discussion

3:30 – 4:00: Coffee break

Session III
4:00 – 4:15:

"Information visualization: the broad product design problem"
Steve Roth, CMU

4:45 – 5:30: Discussion for NSF report

Conference on Automated Learning and Discovery

Workshop W6, Friday, June 12th
Rangos 1, University Center

Learning from Text and the Web

Chaired by: Jaime Carbonell, Steve Fienberg, Tom Mitchell and Yiming Yang (chair)

Session I
11:00 - 11:15: Introduction

11:15 - 12:15: Paper presentations (15 min each)
"Wrapping Semistructured Web Pages with Finite-State Transducers"
Chun-Nan Hsu, Ming-Tzung Dung

"Building Intelligent Agents for Web-Based Tasks:A Theory-Refinement Approach"
Jude Shavlik, Tina Eliassi-Rad

"Learning to Exploit Document Relationships and Structure: The Case for Relational Learning on the Web"
Sean Slattery, Mark Craven

"Wrapper Induction for Semistructured, Web-based Information Sources"
Ion Muslea, Steve Minton, Craig Knoblock
12:15 - 12:30: Discussion

12:30 – 2:00: Lunch

Session II
2:00 - 2:45: Paper presentations (15 min each)
"CiteSeer: An Autonomous System for Processing and Organizing Scientific Literature on the Web"
Kurt Bollacker, Steve Lawrence, C. Lee Giles

"Simultaneous Word and Document Clustering"
John Lafferty, Peter Venable

"Learning and Representing Topic, A Hierarchical Mixture Model for Word Occurrences in Document Databases"
Thomas Hofmann
2:45 - 2:55: Discussion
2:55 - 3:40: Paper presentations (15 min each)
"Event Detection"
Yiming Yang, Tom Pierce, Jaime Carbonell

"Text Mining with Decision Rules and Decision Trees"
Chid Apte, Fred Damerau, Sholom Weiss

"Feature selection for classification based on text hierarchy"
Dunja Mladenic, Marko Grobelnik

3:40 - 3:55 Discussion

3:55 – 4:25: Break

Session III
4:20 -5:20:
Panel
Jaime Carbonell, Ray Mooney, Jay Pedersen, Oren Etzioni, William Cohen, James Allan

Conference on Automated Learning and Discovery

Workshop W7, Friday, June 12th
Rangos 2, University Center

Robot Exploration and Learning

Chaired by: Howie Choset, Maja Matarić and Sebastian Thrun

Session I – What is the role of statistics and reinforcement learning in robot learning?

11:00-11:20: Invited Talk
"Why Robbie Can't Learn"
Leslie Kaelbling

11:20-11:30: Discussion

11:30-12:10: Short presentations (2 slides each)

- Team-Partitioned, Opaque-Transition Reinforcement Learning *(by Peter Stone and Manuela Veloso)*
- A Learning Algorithm for Cooperating Autonomous Agents *(by Christopher Zach)*
- Learning Approaches in Cooperative Robotics: A Few Facts *(by Claude Touzet)*
- Reinforcement Learning of 3-D Object Recognition from Appearance *(by L. Paletta, M. Prantl, and A. Pinz)*
- Robotic Exploration Via Partially Observable Markov Decision Processes *(by Michael Littman)*
- Learning Exploration Polices with Models *(by Marco Wiering and Juergen Schmidhuber)*
- Learning Stationary Temporal Probabilistic Networks *(by Daniel Nikovski)*

12:10-12:30: Discussion

12:30-2pm: Lunch

Session II – What is the the role of biology/psychology in robot learning?

2:00-2:20: Invited talk:
"Biological Inspirations for Facilitating Reinforcement Learning in Challenging Domains"
Maja Matarić

2:20-2:30: Discussion

2:30-3:10: Short presentations (2 slides each)

- Socially Embedded Learning: A Learning Framework using Natural Communication between Robots and Humans *(by Hideki Asoh and Toshihiro Matsui)*
- Comparison of Reinforcement Learning Techniques for Automatic Behavior Programming *(by J. Andrew Bagnell, Keith Doty, and A. Antonio Arroyo)*
- The Robot With Seven Brains: Modularity and Re-Use in Reinforcement Learning *(by Malcolm Ryan)*
- Contextual Control Policy Selection *(by Jefferson Coelho Jr., Elizeth G. Araujo, Manfred Huber, and Roderic A. Grupen)*
- Innate Behavior can Train Supervised Neural Network Controllers *(by Noel Sharkey)*
- Efficient representation of environments from collection of images *(by Ben Krose)*

3:10-3:30: Discussion

3:30-4:00: Coffee break

Session III – What key technologies are required to enable realistic implementation?

4:00-4:20: Invited talk:
"Path planning for unknown, partially-known, and changing environments"
Tony Stentz

4:20-4:30: Discussion

4:30-5:00: Short presentations (2 slides each)

- Getting Reinforcement Learning to Work on Real Robots *(by William Smart and Leslie Kaelbling)*
- Learning, Hierarchies, Abstraction and Exploration in Robotics *(by Bruce Digney)*
- Exploration and Spatial Learning Research at NCARAI *(by Brian Yamauchi, Alan Schultz, William Adams, and Kevin Graves)*
- Acquiring maps with mobile robots using Generalized Voronoi Graphs *(by Howie Choset)*
- Distributed Map-making Using Online Generalized Voronoi Graph *(by James Jennings, Chris Kirkwood-Watts and Craig Tanis)*
- On-board Science Understanding: NASA Ames's Efforts *(by Ted L. Roush, Peter Cheeseman, Virginia Gulick, David Wolf, Paul Gazis, Gretchen Benedix, Wray Buntime, Clark Glymor, Liam Pedersen, Mark Ruzon, and Jon Oliver)*

5:00-5:30: Discussion and Summary

Conference on Automated Learning and Discovery

Workshop W8, Friday, June 12th
Dowd Room, University Center

The Neural Basis of Learning

Chaired by: David Touretzky, Tai Sing Lee and Bill Skaggs

Much has been discovered about the neural basis of learning in humans and animals. For this workshop we will survey three areas where interesting connections might be made between neuroscience and machine learning. As the audience is not expected to have a biological background, each session will begin with some tutorial information. But we hope to rapidly move from tutorial format to an informal discussion about the current state of understanding in each area, and how insights from machine learning can be guide the development of neuroscience theory.

Areas to be covered:

- Dopamine as a reward or novelty signal: a possible mechanism for reinforcement learning in the brain.
- Models of learning and development in visual cortex.
- Animal learning: current theories of classical and instrumental conditioning.