December 12, 2008, 07:30–10:40 and 15:30–18:40 Hilton: Black Tusk **WS7**

Machine Learning Meets Human Learning

http://pages.cs.wisc.edu/~jerryzhu/nips08.html

Nathaniel Daw daw@gatsby.ucl.ac.uk
New York University
Tom Griffiths tom_griffiths@berkeley.edu
University of California, Berkeley
Josh Tenenbaum jbt@mit.edu
Massachusetts Institute of Technology
Xiaojin Zhu jerryzhu@cs.wisc.edu
University of Wisconsin-Madison

Abstract

Can statistical machine learning theories and algorithms help explain human learning? Broadly speaking, machine learning studies the fundamental laws that govern all learning processes, including both artificial systems (e.g., computers) and natural systems (e.g., humans). It has long been understood that theories and algorithms from machine learning are relevant to understanding aspects of human learning. Human cognition also carries potential lessons for machine learning research, since people still learn languages, concepts, and causal relationships from far less data than any automated system. There is a rich opportunity to develop a general theory of learning which covers both machines and humans, with the potential to deepen our understanding of human cognition and to take insights from human learning to improve machine learning systems. The goal of this workshop is to bring together the different communities that study machine learning, cognitive science, neuroscience and educational science. We will investigate the value of advanced machine learning theories and algorithms as computational models for certain human learning behaviors, including, but not limited to, the role of prior knowledge, learning from labeled and unlabeled data, learning from active queries, and so on. We also wish to explore the insights from the cognitive study of human learning to inspire novel machine learning theories and algorithms. It is our hope that the NIPS workshop will provide a venue for cross-pollination of machine learning approaches and cognitive theories of learning to spur further advances in both areas.

7.30-7.35	**Welcome** Workshop Organizers
7.35-7.55	**Training Deep Architectures: Inspiration from Humans** Yoshua Bengio
7.55-8.15	**Stochastic programs as a framework for clustering and causation** Noah Goodman
8.15-8.35	**Learning abstract causal knowledge: a case study in human and machine learning** Josh Tenenbaum
8.35-8.55	**Rational Approximations to Rational Models of Categorization** Adam Sanborn
8.55-9.10	Coffee
9.10-9.30	**The role of prior knowledge in human reconstructive memory** Mark Steyvers

9.30-9.50	**Normative models of multiple interacting memory systems** Mate Lengyel
9.50-10.10	**Compositional Logic Learning** Alan Yuille
10.10-10.40	Panel Discussion 1: Probabilistic models of cognition
10.40-	Poster or Ski
15.30-15.50	**Reconciling reinforcement learning and risk sensitivity: a model-based fMRI study** Yael Niv
15.50-16.10	**Goal-directed decision making as structured probabilistic inference** Matthew Botvinick
16.10-16.30	**Reward bonuses for efficient, effective exploration** Michael Littman
16.30-16.50	**Human Semi-Supervised Learning and Human Active Learning** Xiaojin Zhu
16.50-17.05	Coffee
17.05-17.25	**Where am I and what should I do next? Overcoming perceptual aliasing in sequential tasks** Todd Gureckis
17.25-17.45	**Using reinforcement learning models to interpret human performance on Markov decision problems** Michael Mozer
17.45-18.05	**A Bayesian Algorithm for Change Detection with Identification: Rational Analysis and Human Performance** Jun Zhang
18.05-18:40	Panel Discussion 2: Decision and reward
18.40-	Poster

Training Deep Architectures: Inspiration from Humans
Yoshua Bengio, Universite de Montreal
Theoretical results in circuit complexity suggest that deep architectures are necessary to efficiently represent highly-varying functions, which may be needed for many AI tasks. However, training deep architectures is not only non-convex, but the optimization difficulty seems to increase for deeper architectures. Can we get inspiration from how humans manage to learn complicated concepts and high-level abstractions? The first successful algorithms for training deep architectures suggest a principle is at work: first optimizing something easier (learning concepts that can represented with shallower architectures), and gradually increasing the difficulty (increasing depth), in such a way as to guide the optimization towards better basins of attraction of a local optimization procedure. Another related principle that we are exploring involves breaking from the traditional iid dataset methodology, and breaking training into gradually more difficult phases (and data streams), where each phase allows the learner to learn more complex concepts, exploiting previously learned concepts. Other inspiration from how humans learn complicated concepts will be discussed.

Stochastic programs as a framework for clustering and causation
Noah Goodman, MIT

I will consider a series of concept learning problems faced by people in everyday life. These will proceed from simple clustering problems to the problem of learning latent events underlying a stream of input and the causal relations amongst these events. Each learning problem will be formulated as a stochastic program in the Church language, and I will argue that this framework permits flexible and rapid investigation of learning problems for both cognitive science and machine learning.

Learning abstract causal knowledge: a case study in human and machine learning
Josh Tenenbaum, MIT
TBA

Rational Approximations to Rational Models of Categorization
Adam Sanborn, Gatsby, University College London
Rational models have been successfully used to explain behavior as the optimal solution to a computational problem in many areas of cognition, including memory, reasoning, generalization, and causal induction. While these models can be used to explore the assumptions people make in a particular task, the computation required to produce the optimal solution is often intractable and thus not a reasonable model of the computations performed by people. To make working with rational models practical, computer scientists have developed approximation algorithms with asymptotic convergence guarantees, such as Gibbs sampling and particle filtering. We propose to use these same algorithms to generate rational process models from rational models of cognition – making the assumption that cognition utilizes these statistical algorithms to approximate intractable rational models. In particular, we show that a particle filter approximation to the Rational Model of Categorization (RMC; Anderson, 1990) can reproduce human data, including more human-like order effects than are produced by the RMC.

The role of prior knowledge in human reconstructive memory
Mark Steyvers, University of California, Irvine
Prior knowledge and expectations about events are known to influence recall in human memory, but the specific interactions of memory and knowledge are unclear. We propose hierarchical Bayesian models of reconstructive memory in which prior knowledge is combined with noisy memory representations at multiple levels of abstraction. We present empirical evidence from studies where participants reconstruct the sizes of objects, recall objects in scenes and draw handwritten digits from memory. These studies demonstrate the hierarchical influences of prior knowledge and the beneficial effects of utilizing prior knowledge in recall.

Normative models of multiple interacting memory systems
Mate Lengyel, University of Cambridge
In this talk I will demonstrate how ideas from machine learning, namely unsupervised learning, reinforcement learning, and information theory, can be used to understand fundamental aspects of semantic, episodic, and working memory, respectively, and the interaction of these memory systems in particular. We developed a normative theory of learning about meaningful chunks in visual scenes, and of the way such statistically optimal representations on long-term memory should affect short-term retention of visual scenes in working memory. We also investigated why and how even such a seemingly optimal system might still be beaten by a much simpler episodic memory-based system when one considers the ultimate use of memories for decision making. Most of the work I will present also includes experimental data, collected by collaborators, that test key predictions of the theories.

Compositional Logic Learning
Alan Yuille, UCLA
The paper describes a new method for learning conditional probabilities from binary-valued labeled data. We represent the distributions in noisy-logical form (Yuille and Lu 2008) which is motivated by experiments in Cognitive Science and which offers an alternative to the sigmoid regression representation used (implicitly) by methods like AdaBoost. We specify algorithms for learning these distributions by composing them from elementary structures. Our experimental results show that we obtain experimental results which are slightly better than AdaBoost but which are of far simpler forms.

Reconciling reinforcement learning and risk sensitivity: a model-based fMRI study

Yael Niv, Princeton

Which of these would you prefer: getting $10 with certainty or tossing a coin for a 50% chance to win $20? Whatever your answer, you probably were not indifferent between these two options. In general, human choice behavior is influenced not only by the expected reward value of options, but also by their variance, with subjects differing in the degree to which they are risk-averse or risk-seeking. Traditional reinforcement learning (RL) models of action selection, however, rely on temporal difference methods that learn the mean value of an option, ignoring risk. These models have been strongly linked to learning via prediction errors conveyed by dopaminergic neurons, and to BOLD signals reflecting prediction errors in the nucleus accumbens. Here, in an fMRI study of decision making, we set forth to reconcile the behavioral results and computational theory by inquiring whether the neural implementation of RL is indeed risk-neutral or whether it shows sensitivity to risk. We used the neural signature of RL in the nucleus accumbens to compare between four qualitatively different computational models of how risk can influence decision making. Our results reveal that choice behavior is better accounted for by incorporating risk-sensitivity into reinforcement learning, and, furthermore, that the BOLD correlates of prediction error learning in the brain indeed reflect subjective risk-sensitivity.

Goal-directed decision making as structured probabilistic inference

Matthew Botvinick, Princeton University

Within psychology and neuroscience, there is growing interest in the mechanisms underlying "goal-directed" decision making: the selection of actions based on 1) knowledge of action-outcome contingencies, and 2) knowledge of the incentive value associated with specific outcomes. In formulating theories of how humans and other animals accomplish this kind of decision making, it is natural to look to classical methods for solving Markov decision problems. However, some additional leverage may be gained by considering a more recent approach, which translates the dynamic programming task into a problem of structured probabilistic inference. I'll describe one version of this approach, involving recursive Bayesian inference within graphical models. The components of the underlying graphs align with a set of key functional anatomical systems, allowing the theory to make contact with cognitive neuroscientific data. The approach also gives rise to novel predictions concerning human choice behavior, some of which we have been testing through experimental work.

Reward bonuses for efficient, effective exploration

Michael Littman, Rutgers University

Children must strike a balance between taking the time to perfectly understand their environment and taking advantage of what they already know. Viewed mathematically, solving this exploration-exploitation dilemma is computationally difficult, even in the case in which the environment simply consists of two unknown values (so-called 'bandit' problems). Natural environments present an even more challenging problem because the number of possible events to consider learning about vastly outnumbers the opportunities to explore. In practice, children can never completely experience their world, but nonetheless need to understand it well enough to navigate, make predictions, and explain the events around them.

The exploration-exploitation dilemma has long been recognized in the engineering disciplines as a problem that learning systems must face. Recent work in computer science has highlighted the importance of re-treating from perfect optimality and settling for 'good enough' solutions. This talk will survey some new developments in machine learning that introduce reward bonuses for insufficiently explored states and show that the resulting learning algorithms balance exploration and exploitation while remaining computationally tractable. They can also search hypotheses spaces, even given noisy experience, to find rules that allow them to make predictions in the absence of exhaustive experience. These computationally tractable solutions from the machine-learning community could provide insight on the potential limitations, constraints, and mechanisms that may shape children's exploration and understanding of the world.

Human Semi-Supervised Learning and Human Active Learning

Xiaojin Zhu, University of Wisconsin-Madison

We explore the connections between machine learning and human learning in two settings: semi-supervised learning and active learning. Both are well studied in statistical machine learning. In our experiments, humans replace learning algorithms to assume the role of the learner in a category learning (i.e., classification) task. In semi-supervised learning, subjects are given additional unlabeled data. In active learning, subjects are given the ability to choose which items to query for label. Our results indicate that humans can perform semi-supervised learning and active learning. Quantitatively their performance also differs from learning theory predictions in interesting ways.

Where am I and what should I do next? Overcoming perceptual aliasing in sequential tasks
Todd Gureckis, New York University
A critical challenge facing learners in a changing environment is correctly representing the current state of the world and appreciating how it may influence future outcomes. My talk considers recent work in my lab looking at issues of state representation and generalization in sequential decision making by humans. The experiments and models I describe are principally inspired by recent advances in machine learning which address how artificial agents may learn from experience in complex task domains. Overall, the goal of this work is to establish connections between this foundational computational work and issues of mental representation, categorization, stimulus generalization, and decision making traditionally studied in cognitive science/psychology.

Using reinforcement learning models to interpret human performance on Markov decision problems
Michael Mozer, University of Colorado at Boulder
Theories of learning by reinforcement have been used to interpret data from individuals performing one-step choice tasks (e.g., the Iowa gambling task), and data from animals performing temporally extended behaviors, but not, to our knowledge, data from individuals performing sequential decision tasks. We tested participants in a temporally extended task that involved exploring an unfamiliar environment. The environment consisted of rooms, each containing two doors leading to other rooms. The participant's task was to select a sequence of doors to enter. Rewards were associated with state-action pairs. One question we address is whether formal theories of reinforcement learning (Q learning, Q policy gradient, and model based approaches) are suitable for characterizing the behavior of participants. We obtained a maximum likelihood fit of Q learning parameters to the pattern of choices made by individual participants. The parameters include: exploration strategy (epsilon-greedy versus normalized exponential), control of the exploration-exploitation trade off, the discounting rate (gamma), the backup parameter of the eligibility trace (lambda), and a learning rate. We report mixed results fitting participant data to the models. Beyond using reinforcement-learning models to fit data, the data has the potential to inform theories of reinforcement learning. These theories are neutral with regard to how the model parameters are set. Thus, a second question we address is: how do task variables and cognitive constraints modulate parameter settings? To explore this question, we performed experimental manipulations such as varying the time allotted for choosing an action, and varying a concurrent working-memory load. We find that these manipulations can be interpreted in terms of their influence on model parameters.

A Bayesian Algorithm for Change Detection with Identification: Rational Analysis and Human Performance
Jun Zhang, University of Michigan-Ann Arbor
We consider the problem of change detection along with identification in multi-hypotheses setting, where the state-of-world changes from H_0 to H_i $(i = 1, 2, \ldots, N)$ under known prior distributions. A Bayesian sequential updating equation is derived, along with the usual boundary-crossing stopping rule. The algorithm has the property that the value of an absorbing boundary, when overshoot is ignored, equals the hit rate of a decision-maker conditioned on that response. Computer simulation reveals that the algorithm shares many similarities with human performance in stimulus detection/identification experiments.

POSTERS

1. A psychophysical investigation of clustering. Joshua Lewis, UCSD

2. The Hierarchical Dirichlet Process as a model of Human Categorization. Kevin Canini, Berkeley

3. Kernels and Exemplar Models. Frank Jäkel, MIT
4. Learning Object-based Attention Control. Ali Borji, Majid N. Ahmadabadi and Babak N. Araabi, Institute for Studies in Theoretical Physics and Mathematics, Iran
5. A Hebbian Learning Rule for Optimal Decision Making. Michael Pfeiffer, Bernhard Nessler, and Wolfgang Maass, Graz University of Technology, Austria
6. Modeling Word Association Data using Multiple Maps. Laurens van der Maaten and Geoffrey Hinton, Tilburg University and University of Toronto
7. Integrating Statistics from the World and from Language to Learn Semantic Representations. Mark Andrews, Gabriella Vigliocco, and David P. Vinson, University College London
8. Bayesian modeling of intuitive pedagogical reasoning. Patrick Shafto and Noah Goodman, University of Louisville.
9. Learning from actions and their consequences: Inferring causal variables from continuous sequences of human action. Daphna Buchsbaum and Tom Griffiths, University of California Berkeley
10. Translation-invariant sparse deep belief networks for scalable unsupervised learning of hierarchical representation. Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Y. Ng, Stanford University
11. Machine learning in the service of understanding human learning: an ideal observer-based analysis of the learning curve. Ferenc Huszar, Uta Noppeney, and Mate Lengyel. Budapest University of Technology and Economics, MPI Tuebingen, and University of Cambridge