Value of Information in Inference, Learning and Decision-Making

Saturday December 10, 2005

Alina Beygelzimer, Rajarshi Das, Irina Rish (primary contact), Gerry Tesauro
IBM T. J. Watson Research Center

www.research.ibm.com/nips05workshop/

Overview and Goals A common fundamental problem of value of information (VOI) analysis arises in inference, learning and sequential decision-making when one is allowed to actively select, rather than passively observe, the input information. VOI provides a principled methodology that enables acquiring information in a way that optimally trades off the cost of information gathering with the expected benefit in some overall objective (e.g., classification accuracy or cumulative reward).

For example, in Bayesian problem diagnosis VOI analysis aims at selecting observations (e.g., medical tests) that are most informative about the unknown variables (e.g., diseases we are trying to diagnose) while minimizing the cost of collecting the information. In sequential decision-making problems, VOI can provide a principled solution to the well-known "exploration versus exploitation" dilemma, so that one can optimally trade off the immediate cost of exploratory actions with expected improvement in future decisions and future reward. Yet another example is active learning, where the goal is to minimize the cost of observations (e.g., the number of labeled samples) while maximizing the learner's objective function. Finally, selecting the most relevant subset of features in supervised learning is another example where VOI analysis can provide a principled solution.

Clearly, these areas differ in their choices of a particular objective function and the approaches to active exploration, but have a common goal of selecting explorative actions that maximize the VOI. In this workshop, we plan to bring together researchers from several fields concerned with VOI analysis and hope to ignite cross-fertilization between the areas. This could lead to major theoretical progress as well as practical impact in applications such as medical diagnosis, quality control in product design, IT systems management and troubleshooting, and DNA library screening, just to name a few.

Value of Information in Inference, Learning and Decision-Making

Saturday December 10, 2005

Organizers: Alina Beygelzimer, Rajarshi Das, Irina Rish, Gerry Tesauro

Morning session: 7:30am–10:30am

7:30am **Introduction**, *Irina Rish*

7:45am **Value of Information in Interactive Optimization**, *Craig Boutilier*

8:15am **Nonmyopic Value of Information in Graphical Models**, *Carlos Guestrin*

8:45am *Break and Poster Setup*

9:00am **Models of Active Learning**, *Sanjoy Dasgupta*

9:30am *Poster Spotlights*

10:20am *Poster Session*

Afternoon session: 3:30pm–6:30pm

3:30pm **PAC bounds on the exploration-exploitation tradeoff**, *Michael Littman*

4:00pm **Bayesian Sparse Sampling for On-line Reward Optimization**, *Dale Schuurmans*

4:30pm Optimal Experimental Design Models of Naive Human Information Acquisition, *Jonathan Nelson, Gary Cottrell, Flavia Filimon, and Terry Sejnowski*

4:45pm *coffee break*

5:00pm Compressed Sensing and Active Learning, *Rui Castro, Jarvis Haupt, and Robert Nowak*

5:15pm Active Learning in POMDPs, *Robin Jaulmes, Joelle Pineau, and Doina Precup*

5:30pm When Ignorance is Bliss, *Peter Grunwald and Joseph Halpern*

5:45pm *Concluding Discussion*

Abstracts of Invited Talks

Speaker: Craig Boutilier, University of Toronto, Canada (**7:45am**)
Title: *Value of Information in Interactive Optimization*
Abstract Interactive optimization requires assessing a decision maker's preferences (or an objective function) during the course of optimization (e.g., recommending a decision or sequence of decisions to the decision maker). When interactions costs such as time costs or cognitive burden are considered, the value of information contained in any interaction proposed by a decision-support system must be traded off against these costs. In this talk, I will describe a general framework for interactive optimization and decision support in which the role of value of information is central. This framework will reveal the the close ties concepts in reinforcement learning, exploration vs. exploitation, and revealed preference (or inverse RL).

Speaker: Carlos Guestrin, Carnegie Mellon University (**8:15am**)
Title: *Nonmyopic Value of Information in Graphical Models*
Abstract In decision making under uncertainty, where one can choose among several expensive queries, it is a central issue to decide which variables to observe in order to achieve a most effective increase in expected utility. This problem has previously only been approached myopically, without any known performance guarantees. In this talk, I will present efficient nonmyopic algorithms for selecting an optimal subset of observations and for computing an optimal conditional plan for a class of graphical models containing Hidden Markov Models. I will also show how our methods can be used for interactive structured classification and for sensor scheduling in a Civil Engineering domain. Many graphical models tasks which can be efficiently solved for chains, can be generalized to polytrees. I will present surprising hardness results, showing that the optimization problems are wildly intractable (NP^{PP} complete) even in the case of discrete polytrees. Addressing these theoretical limits, I will present efficient approximation algorithms for selecting informative subsets of variables. Our algorithms are applicable to a large class of graphical models, and provide a constant factor approximation guarantee of 1-1/e, which is provably the best constant factor achievable unless P = NP. I will sketch how our methods can be extended to optimal experimental design in Gaussian processes, and present extensive evaluation of our algorithms on several real-world data sets. The material presented in this talk is joint work with Andreas Krause.

Speaker: Sanjoy Dasgupta, University of California, San Diego (**9:00am**)
Title: *Models of Active Learning*
Abstract I will give an overview of the various models of active learning that have been proposed, including generative models (such as Castelli-Cover), query by committee, learning from membership queries, and some recent developments. I will talk about the strengths and weaknesses of these models, the kinds of learning algorithms they motivate, and the theoretical results that have been shown.

Speaker: Michael Littman, Rutgers University (**3:30pm**)
Title: *PAC bounds on the exploration-exploitation tradeoff*
Abstract Paraphrasing Frost, a system cannot make two decisions and be one decision maker—it can only observe the effects of the actions it actually chooses to make. The resulting exploration-exploitation tradeoff can be solved optimally only for the simplest decision-making scenarios,

necessitating the development of weaker goals. The PAC framework has emerged as a viable alternative, in which the decision maker has a polynomial number of interactions to obtain near optimal reward with high probability. I will survey existing PAC results for k-armed bandit problems and Markov decision processes, as well as progress from my own group in solving related decision problems.

Speaker: Dale Schuurmans, University of Alberta (**4:00pm**)
Title: *Bayesian Sparse Sampling for On-line Reward Optimization*
Abstract I will discuss progress toward developing sparser "sparse sampling" techniques for approximating Bayes optimal decision making in reinforcement learning. Our approach combines sparse sampling with Bayesian exploration to achieve improved decision making while controlling computational cost. The idea is to grow a sparse lookahead tree, intelligently, by exploiting information in a Bayesian posterior—rather than enumerate action branches (standard sparse sampling) or compensating myopically (value of perfect information). The outcome is a flexible, practical technique for improving action selection in simple reinforcement learning scenarios. Joint work with Tao Wang, Dan Lizotte and Mike Bowling.

Contributed Talks

- *Optimal Experimental Design Models of Naive Human Information Acquisition*, Jonathan Nelson, Gary Cottrell, Flavia Filimon, and Terry Sejnowski
- *Compressed Sensing and Active Learning*, Rui Castro, Jarvis Haupt, and Robert Nowak
- *Active Learning in POMDPs*, Robin Jaulmes, Joelle Pineau, and Doina Precup
- *When Ignorance is Bliss*, Peter Grunwald and Joseph Halpern

Posters

- *Nonmyopic VOI with Near-Optimal Collection Cost*, Andreas Krause, Carlos Guestrin, Anupam Gupta, and Jon Kleinberg
- *Efficient VOI for Graphical Models*, Brigham Anderson and Andrew Moore
- *Efficient Selection of Disambiguating Actions for Stereo Vision*, Monika Schaeffer and Ronald Parr
- *Mechanism Design via Machine Learning*, Maria-Florina Balcan, Avrim Blum, Jason Hartline, and Yishay Mansour
- *Gaussian Process Regression for Optimization*, Daniel Lizotte, Tao Wang, Michael Bowling, and Dale Schuurmans
- *Active Learning under Arbitrary Distributions*, Claire Monteleoni and Matti Kääriänen
- *Active Learning, Loss Function Convexity, and Support Vectors*, Leon Bottou, Ronan Collobert, Seyda Ertekin, and Jason Weston
- *Using Value of Information to Learn and Classify under Hard Budgets*, Russell Greiner
- *Active Learning and Optimal Predictions*, Susanne Still
- *Efficient Test Selection in Active Diagnosis via Entropy Approximation*, Alice Zheng, Irina Rish, Alina Beygelzimer