



Call For Papers

NIPS Workshop on



Value of Information in Inference, Learning and Decision-Making

Held at Neural Information Processing Systems (NIPS) conference

Whistler, CANADA: December 10, 2005

Information

Workshop URL

www.research.ibm.com/nips05workshop/

Submission

nips05workshop@watson.ibm.com

NIPS

<http://www.nips.cc>

Dates & Deadlines

October 24: Abstract Submission

October 31: Acceptance Notification

Organizing Committee

Dr. Alina Beygelzimer

IBM T. J. Watson Research Lab, USA

Dr. Rajarshi Das

IBM T. J. Watson Research Lab, USA

Dr. Irina Rish

IBM T. J. Watson Research Lab, USA

Dr. Gerry Tesauro

IBM T. J. Watson Research Lab, USA

Invited Speakers

Prof. Sanjoy Dasgupta

University of California, San Diego

Prof. Carlos Guestrin

Carnegie Mellon University

Prof. Michael Littman

Rutgers University

Prof. Dale Schuurmans

University of Alberta

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.

Suggested Topics

The list of possible topics includes (but is not limited to) the following:

- VOI analysis in probabilistic inference and decision theory
- feature selection and attribute-efficient learning
- active learning (query learning, selective sampling)
- exploration-exploitation trade-off in reinforcement learning
- adaptive versus non-adaptive testing designs
- comparison of different action selection criteria and objective functions
- applications of VOI in diagnosis, systems control and management, coding theory, computational biology, neural coding, etc.

Format

This is a one-day workshop. There will be several invited talks and tutorials (roughly 30-40 minutes each) and shorter contributed talks from researchers in industry and academia, as well as a panel discussion. We will hold a poster session if we receive a sufficient number of good submissions. The workshop is intended to be accessible to the broader NIPS community and to encourage communication between different fields.

Submission Instructions

We invite submissions of extended abstracts (**up to 2 pages, not including bibliography**) for the short contributed talks and/or posters. The submission should present a high-level description of recent or ongoing work related to the topics above. We will explore the possibility of publishing papers based on invited and submitted talks in a special issue of an appropriate journal.

Email submissions to nips05workshop@watson.ibm.com as attachments in **Postscript or PDF**, no later than October 24, 2005.