


NIPS 2007



NEURAL
INFORMATION
PROCESSING
SYSTEMS
CONFERENCE



Workshops

TUTORIALS


December 3, 2007
Hyatt Regency
Vancouver, BC, Canada

CONFERENCE

December 3-6, 2007
Hyatt Regency
Vancouver, BC, Canada

WORKSHOPS

December 7-8, 2007
Westin Resort & Spa
Hilton Resort & Spa
Whistler, BC, Canada

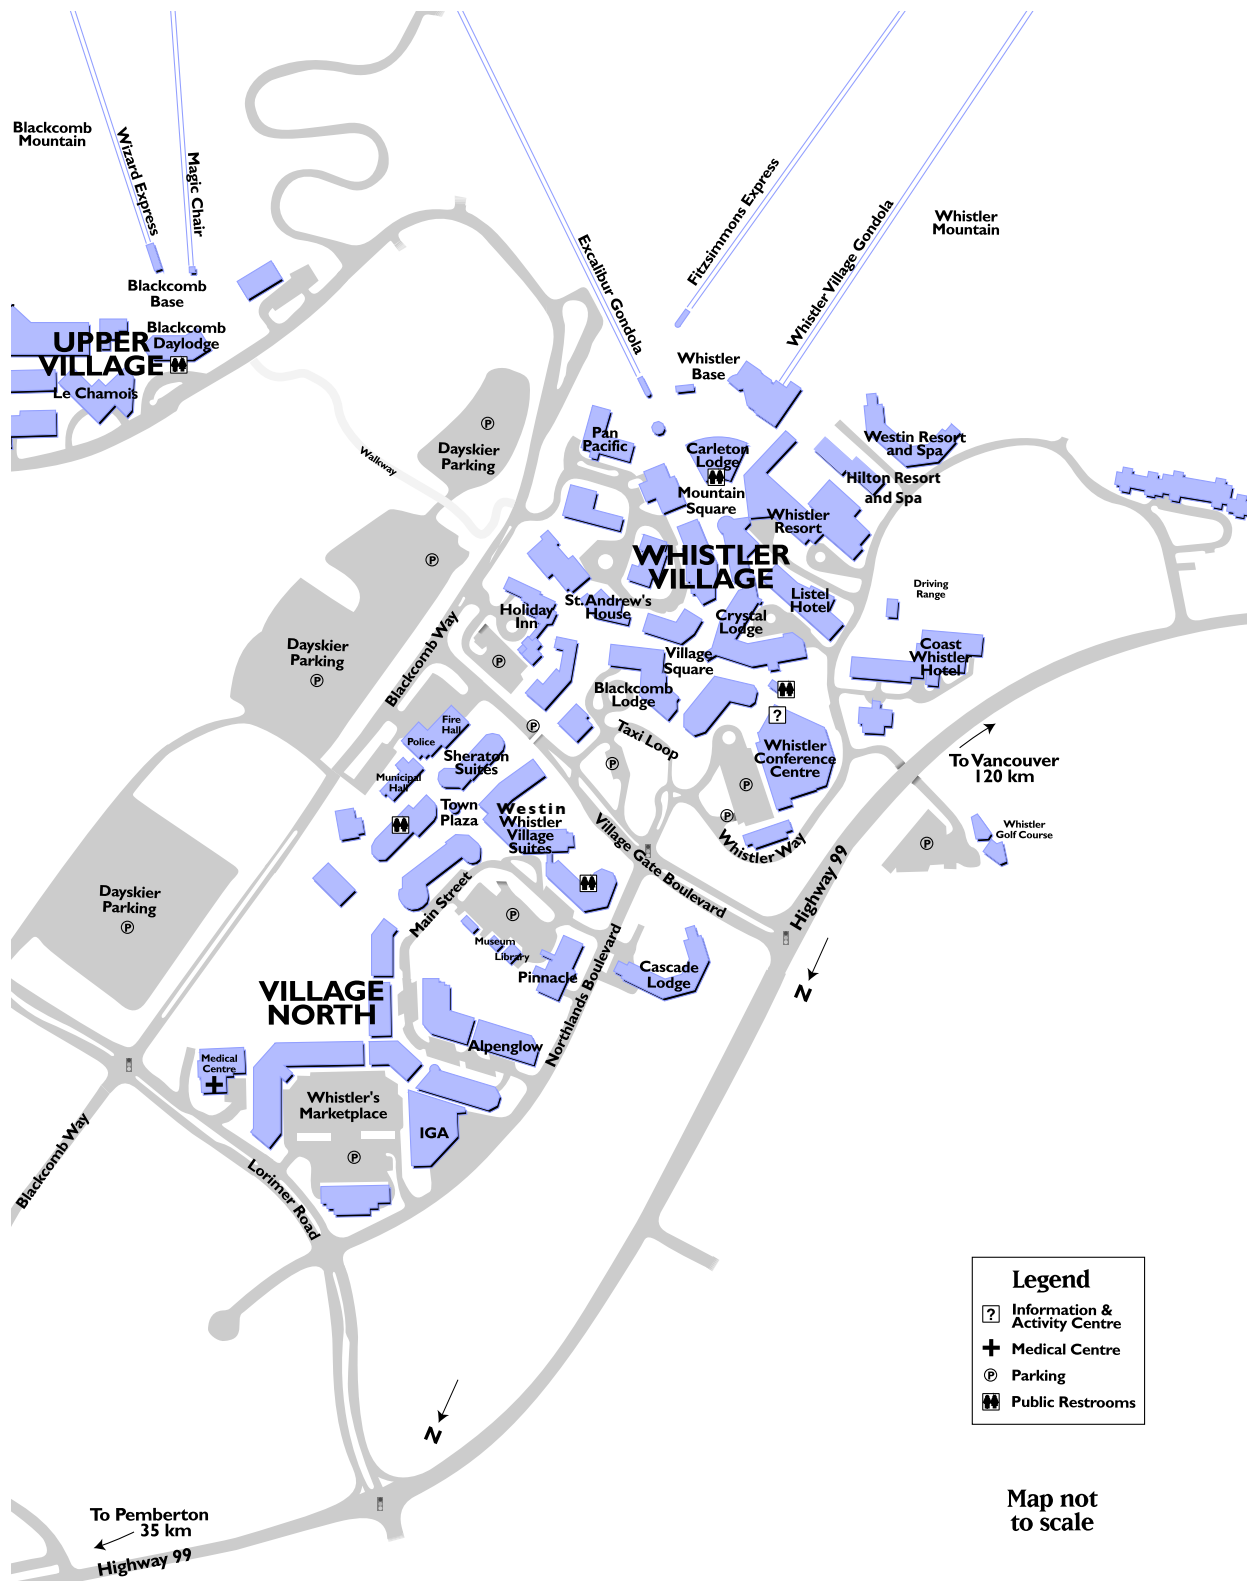


Workshop Program

Westin Resort and Spa
and
Hilton Resort and Spa
Whistler, BC, Canada

Edited by Adrienne L. Fairhall and Robert C. Williamson

Whistler Map



NIPS 2007 Workshops—Program of Events

Thursday, December 6:

2:00–3:30 pm	Buses depart Vancouver Hyatt for Westin Resort and Spa
5:00–9:00 pm	Lift Ticket Sales – West end of Lobby
5:00–8:30 pm	Registration – West end of Emerald Foyer
6:30–8:30 pm	Welcoming Reception – Emerald Ballroom

Friday, December 7:

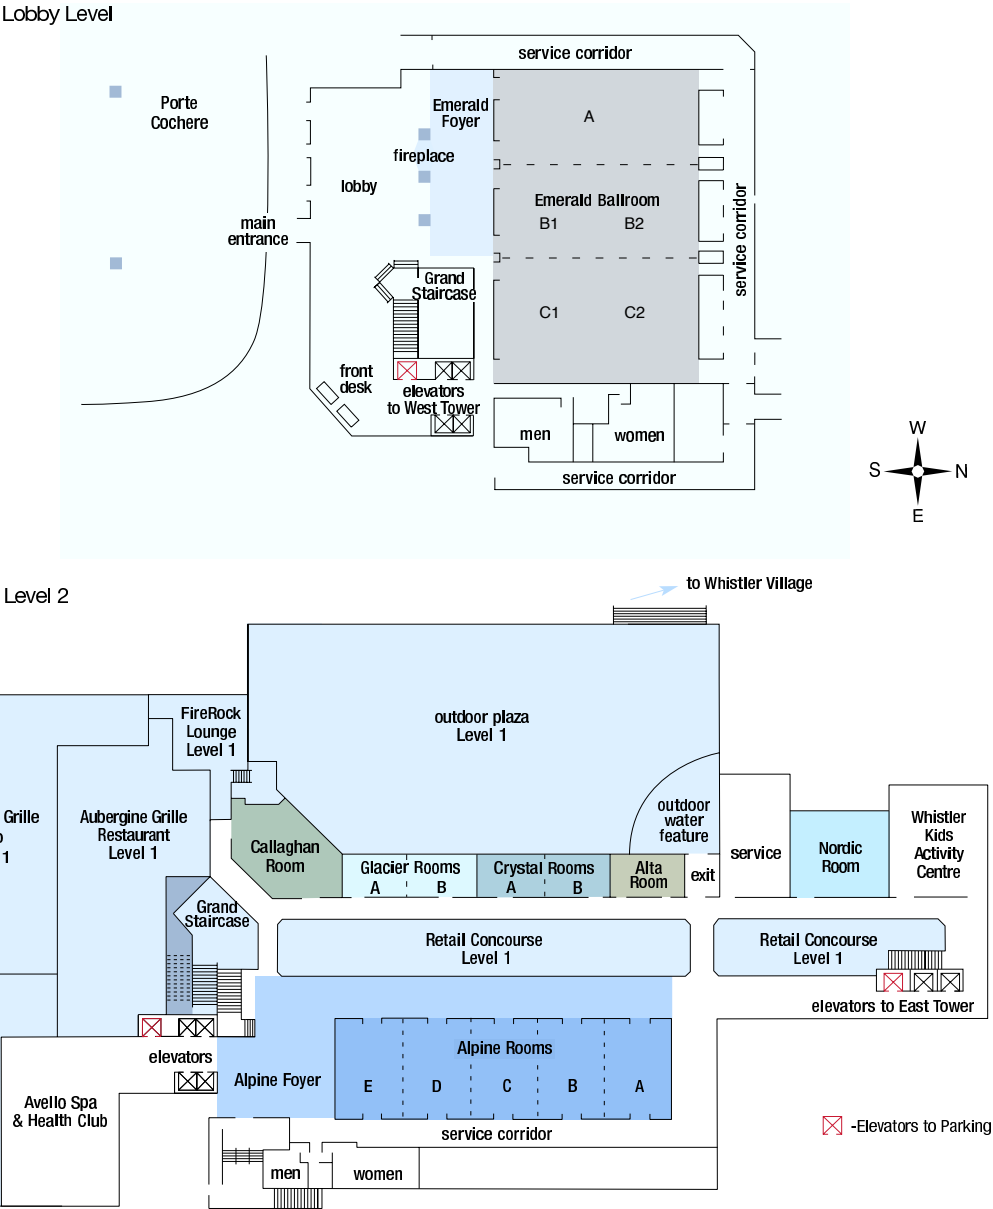
6:30–8:00 am	Breakfast – Westin: Emerald Ballroom
7:00–11:00 am	Registration – West end of Emerald Foyer
7:30–10:30 am	Workshop Sessions*
10:00–noon	Lift Ticket Sales – West end of Lobby
2:00–6:30 pm	Workshop Sessions Continue*

Saturday, December 8:

6:30–8:00 am	Breakfast – Westin: Emerald Ballroom
7:00–11:00 am	Registration – West end of Emerald Foyer
7:30–10:30 am	Workshop Sessions*
10:00–noon	Lift Ticket Sales – West end of Lobby
3:30–6:30 pm	Workshop Sessions Continue*
7:30–10:30 pm	Banquet & Wrap Up Meeting – Emerald Ballroom

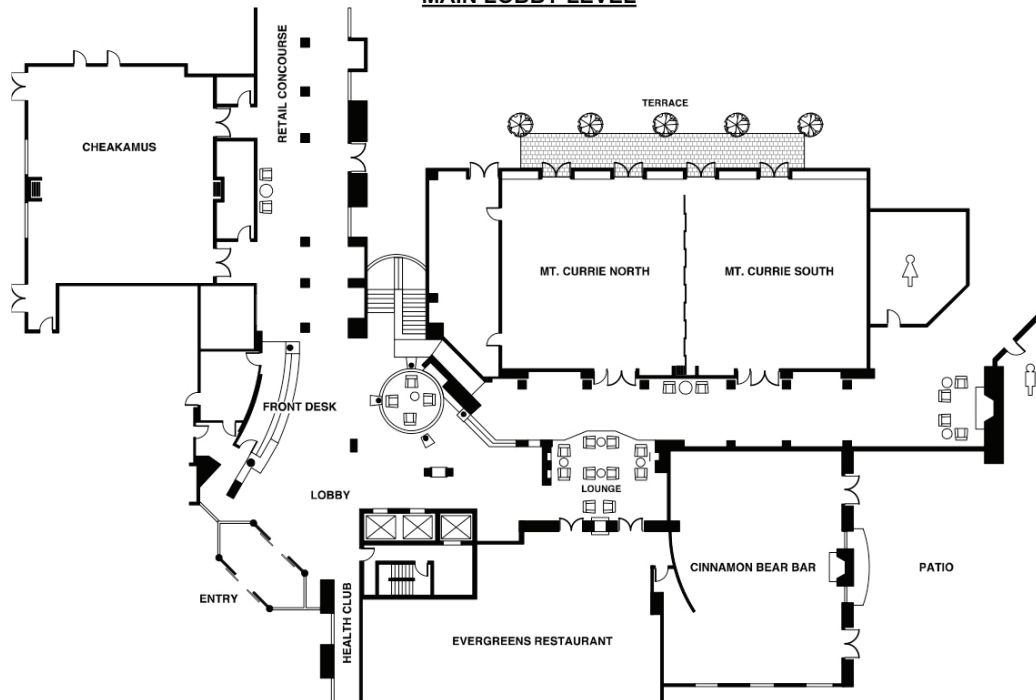
*** Important: A few workshops begin at 9:00 am and end at noon, then start at 2:00pm and end at 5:00pm. Check individual workshop times on the schedule page.**

Westin Resort Workshop Rooms




Hilton Resort Workshop Rooms

Hilton Whistler Resort MAIN LOBBY LEVEL



LOWER LEVEL



Friday
December 7, 2007

Workshops on Friday, December 7th[★]

Principles of Learning Problem Design

John Langford, Alina Beygelzimer

Hilton: Sutcliffe A

Topology Learning: New Challenges at the Crossing of Machine Learning, Computational Geometry and Topology

Michaël Aupetit, Frédéric Chazal, Gilles Gasso, David Cohen-Steiner, Pierre Gaillard

Westin: Glacier

Robotics Challenges for Machine Learning

Jan Peters, Marc Toussaint

Hilton: Black Tusk

Machine Learning for Web Search

Dengyong Zhou, Olivier Chapelle, Thorsten Joachims, Thomas Hofmann

Hilton: Mt Currie N

Approximate Bayesian Inference in Continuous/Hybrid Models

Matthias Seeger, David Barber, Neil Lawrence, Onno Zoeter

Hilton: Diamond Head

Beyond Simple Cells: Probabilistic Models for Visual Cortical Processing

Richard Turner, Pietro Berkes, Maneesh Sahani

Westin: Alpine A–C

TWO DAY WORKSHOPS — DAY 1

Efficient Machine Learning: Overcoming Computational Bottlenecks in Machine Learning (Part 1)

Samy Bengio, Corinna Cortes, Dennis DeCoste, Francois Fleuret, Ramesh Natarajan, Edwin Pednault, Dan

Pelleg, Elad Yom-Tov

Hilton: Mt Currie S

MLSys: Statistical Learning Techniques for Solving Systems Problems (Part 1)

Sumit Basu, Archana Ganapathi, Emre Kiciman, Fei Sha

Hilton: Sutcliffe B

Machine Learning in Computational Biology (Part 1)

Gal Chechik, Christina Leslie, Quaid Morris, William Stafford Noble, Gunnar Rätsch, Koji Tsuda

Westin: Nordic

Large Scale Brain Dynamics (Part 1)

Ryan Canolty, Kai Miller

Westin: Emerald A

Music, Brain and Cognition (Part 1) Learning the Structure of Music and its Effects on the Brain

David R. Hardoon, Eduardo Reck Miranda, John Shawe-Taylor

Westin: Alpine (D–E)

Hierarchical Organization of Behavior: Computational, Psychological & Neural Perspectives (Part 1)

[9:00–12:00; 3:00–6:00]

Yael Niv, Matthew Botvinick, Andrew Barto

Westin: Callaghan

★ All sessions 7:30–10:30 and 3:30–6:30 unless otherwise indicated.

Principles of Learning Problem Design

Friday, December 7, 2007

John Langford, *Yahoo! Research*

Alina Beygelzimer, *IBM Research*

<http://hunch.net/learning-problem-design/>

This workshop is about how to design learning problems. The dominant system for applying machine learning in practice involves a human labeling data. This approach is limited to situations where human experts exist, can be afforded, and are fast enough to solve the relevant problem. In many settings these constraints are not met, yet it appears that machine learning is still possible via cleverly reinterpreting or reusing existing data. The basic idea is to create supervised learning problems from data which is not conventionally labeled in such a way that successfully solving these ancillary problems is helpful in solving the original learning problem. Since the task is akin to the problem of mechanism design in economics and game theory, we call it learning problem design.

Several recent examples of learning problem design include converting otherwise-unsupervised problems into supervised problems; creating recursive prediction problems (predicting from predictions); reducing one learning task to another.

This area is new and not entirely defined. It is our goal to bring together anyone who is interested in the topic, define what we do and don't understand, and attempt to define the principles of learning problem design.

Principles of Learning Problem Design

Friday, December 7, 2007

Organizers: John Langford & Alina Beygelzimer

Morning session: 7:30am–10:30am

- 7:30am **Principles of Learning Problem Design: Introduction**, *John Langford*
- 8:00am **Unsupervised learning as supervised structured prediction**, *Hal Daumé III*
- 8:30am **The Information Theoretic Basis of Multi-View Learning**, *Sham Kakade*
- 9:00am *Coffee Break + Posters*
- 9:15am **Representations in Learning Task Design**, *Mark Reid and Robert Williamson*
- 9:45am **Panel discussion**, *moderated by Alina Beygelzimer*

All morning talks are 20 minutes long and followed by 10 minutes of questions.

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Problem Design in Games**, *Luis von Ahn*
- 4:10pm **Self-Taught Learning**, *Rajat Raina*
- 4:30pm **Generalizing Backpropagation to Incorporate Sparse Coding**, *Dave Bradley*
- 4:50am *Coffee Break + Posters*
- 5:00pm **Spotlights (alphabetically)**: Nina Balcan, Greg Druck, Rakesh Gupta, Huzefa Rangwala, Ashutosh Saxena
- 5:30pm **Panel Discussion and Closing Remarks**, *moderated by John Langford*
- 6:00pm **Poster Time**

Topology Learning New Challenges at the Crossing of Machine Learning, Computational Geometry and Topology

Friday December 7, 2007

Michaël Aupetit, *CEA DAM, Bruyères-le-Châtel, France*

Frédéric Chazal, *INRIA Futurs, France*

Gilles Gasso, *LITIS, INSA Rouen, France*

David Cohen-Steiner, *INRIA Sophia Antipolis, France*

Pierre Gaillard, *CEA DAM, Bruyères-le-Châtel, France*

<http://topolearnnips2007.insa-rouen.fr/index.html>

There is a growing interest in Machine Learning, in applying geometrical and topological tools to high-dimensional data analysis and processing. Considering a finite set of points in a high-dimensional space, the approaches developed in the field of Topology Learning intend to learn, explore and exploit the topology of the shapes (topological invariants such as the intrinsic dimension or the Betti numbers), manifolds or not, from which these points are supposed to be drawn. Applications likely to benefit from these topological characteristics have been identified in the field of Exploratory Data Analysis, Pattern Recognition, Process Control, Semi-Supervised Learning, Manifold Learning and Clustering.

Some works concerning manifolds have been done by researchers from Machine Learning. In the framework of Manifold Learning, nonlinear dimension reduction techniques are used to represent or visualize high-dimensional data or as pre-processing step for supervised or unsupervised learning tasks. However, the final dimension of the output data and the topology of the space they live in, are constrained a priori. In Semi-Supervised Learning, it is intended to perform manifold regularization by taking into account the topology of the shapes using the Laplacian of some proximity graph of the data. A similar approach is also used in Spectral Clustering. But in both cases, the choice of the proximity graph greatly impacts the results, making these methods sensitive to noise or outliers, and there is no objective way to choose a good graph. In the field of Computational Geometry, lots of efforts have been realized in processing geometric data in dimension 2 or 3. Using concepts such as epsilon-samples and restricted Delaunay triangulations, efficient tools for estimating topological and geometrical properties of shapes have been developed. However, it is necessary to assume that the unknown shape is a smooth manifold and that the sample is not too noisy to ensure the correctness of the reconstruction algorithms.

Topology Learning aims at going beyond all these limits by providing methods able to extract geodesic paths or topological invariants from noisy high-dimensional out-of-control sampled shapes with theoretical guarantees and algorithmic tractability, and with the least a priori constraints or knowledge. It appears that the integration in the Machine Learning and Statistics frameworks of the problems we are faced with in Topology Learning, is still in its infancy. So we wish this workshop to ignite cross-fertilization between Machine Learning, Computational Geometry and Topology, likely to benefit to all of them by leading to new approaches, deeper understanding, and stronger theoretical results of the problems carried by Topology Learning.

Topology Learning

New Challenges at the Crossing of Machine Learning, Computational Geometry and Topology

Friday December 7, 2007

Organizers: Michaël Aupetit, Frédéric Chazal, Gilles Gasso,
David Cohen-Steiner & Pierre Gaillard

Morning session: 7:30am–10:30am

7:30am **Welcome**, *M. Aupetit, F. Chazal, G. Gasso, D. Cohen-Steiner, P. Gaillard*

7:35am **Invited talk: Topology and Statistics**, *Pr. P. Niyogi*

8:20am **Invited talk: Manifold Denoising**, *Pr. M. Hein*

9:05am *coffee break and poster session 1*

- **Discrete Laplace Operator on Meshed Surfaces**, *M. Belkin, J. Sun, Y. Wang*
- **Mixed Dimensions Estimation and Clustering in High Dimensional Noisy Point Clouds**, *G. Haro, G. Randall, G. Sapiro*
- **Topologically-Constrained Latent Variable Models**, *R. Urtasun, D. Fleet, T. Darrell, N.D. Lawrence*
- **Kernels, Margins, and Low-dimensional Mappings**, *M.-F. Balcany, A. Blum, S. Vempala*
- **Sampling and Topological Inference for General Shapes**, *F. Chazal, D. Cohen-Steiner, A. Lieutier*
- **Toward Manifold-Adaptive Learning**, *A. M. Farahmand, J.-Y. Audibert*

9:45am **Invited talk: Bregman and Voronoï**, *Pr. J.-D. Boissonnat*

Afternoon session: 3:30pm–6:30pm

3:30pm **Invited talk: Topological Persistence**, *Pr. H. Edelsbrunner*

4:15pm *coffee break and poster session 2*

- **Manifold-based Approaches for Improved Classification**, *M. A. Davenport, C. Hegde, M. B. Wakin and R. G. Baraniuk*
- **Learning the topology of a labeled data set**, *P. Gaillard, M. Aupetit, G. Govaert*
- **Computational Geometry Neighborhoods for Local Learning**, *M. R. Gupta, E. K. Garcia, Y. Chen*
- **Mapper : A Topological Mapping Tool for Point Cloud Data**, *G. Singh, F. Memoli, G. Carlsson*
- **Ultra-Low-Dimensional Embeddings for Doubling Metrics**, *T.-H. Hubert Chan, A. Gupta, K. Talwar*
- **A Probabilistic Perspective on Persistence Homologies**, *SAMSI 2007 working group*
- **Size of Delaunay Triangulation for Points Distributed over Lower-dimensional Polyhedra: a Tight Bound**, *N. Amenta, D. Attali, O. Devillers*

5:10pm **Invited talk: Witness Complexes**, *Dr. V. de Silva*

5:55pm **Panel discussion: Future trends in Topology Learning**,
P. Niyogi, M. Hein, J.-D. Boissonnat, H. Edelsbrunner & V. de Silva

6:25pm **Closing remarks**, *M. Aupetit*

Robotics Challenges for Machine Learning

Friday, December 7, 2007

Jan Peters, *Max-Planck-Institute for Biological Cybernetics*

Marc Toussaint, *Technical University of Berlin*

<http://www.robot-learning.de>

Robotics challenges can inspire and motivate new Machine Learning research as well as being an interesting field of application of standard ML techniques. Despite the wide range of machine learning problems encountered in robotics, the main bottleneck towards autonomous robots has been a lack of interaction between the core robotics and the machine learning communities. To date, many roboticists still discard machine learning approaches as generally inapplicable or inferior to classical, hand-crafted solutions. Similarly, machine learning researchers do not yet acknowledge that robotics can play the same role for machine learning which for instance physics had for mathematics: as a major application as well as a driving force for new ideas, algorithms and approaches.

With the current rise of real, physical humanoid robots in robotics research labs around the globe, the need for machine learning in robotics has grown significantly. Only if machine learning can succeed at making robots fully adaptive, it is likely that we will be able to take real robots out of the research labs into real, human inhabited environments. Among the important problems hidden in these steps are problems which can be understood from the robotics and the machine learning point of view including perceptuo-action coupling, imitation learning, movement decomposition, probabilistic planning problems, motor primitive learning, reinforcement learning, model learning and motor control.

Robotics Challenges for Machine Learning

Friday, December 7, 2007

Organizers: Jan Peters & Marc Toussaint

Morning session: 7:30am–10:30am

- 7:30am **Welcome**, Jan Peters, Max Planck Institute, Marc Toussaint, Technical University of Berlin
- 7:35am **Learning Nonparametric Policies by Imitation**,
David Grimes and Rajesh Rao, University of Washington
- 8:05am **Machine Learning for Developmental Robotics**,
Manuel Lopes, Luis Montesano, Francisco Melo, Instituto Superior Técnico
- 8:15am **Machine Learning Application to Robotics and Human-Robot Interaction**,
Aude Billard, EPFL
- 8:45am *coffee break*
- 9:00am **Poster Spotlights**
- 9:20am **Bayesian Reinforcement Learning in Continuous POMDPs with Application to Robot Navigation**, Joelle Pineau, McGill University
- 9:50am **Self-Supervised Learning from High-Dimensional Data for Autonomous Offroad Driving**, Ayse Naz Erkan, Raia Hadsell, Pierre Sermanet, Koray Kavukcuoglu, Marc-Aurelio Ranzato, Urs Muller, Yann LeCun, NYU
- 10:00am **Task-based Motion Primitives for the Control and Analysis of Anthropomorphic Systems**, Luis Sentis, Stanford University

Afternoon session: 3:30pm–6:30pm

- 3:30am **STAIR: The STanford Artificial Intelligence Robot project**,
Andrew Ng, Stanford University
- 4:00am **Robot Perception Challenges for Machine Learning**,
Chieh-Chih Wang, National Taiwan University
- 4:10am **Probabilistic Inference Methods for Nonlinear, Non-Gaussian, Hybrid Control**,
Nando de Freitas, University of British Columbia
- 4:40pm *coffee break*
- 5:00am **A New Mathematical Framework for Optimal Choice of Actions**, Emo Todorov, UCSD
- 5:30pm **Poster Spotlights**
- 5:50pm **Poster Session**

Machine Learning for Web Search

Friday December 7, 2007

Dengyong Zhou, *Microsoft Research*

Olivier Chapelle, *Yahoo! Research*

Thorsten Joachims, *Cornell University*

Thomas Hofmann, *Google Research*

<http://research.microsoft.com/conferences/nips07/websearch.html>

This workshop is intended for people who are interested in both machine learning and web search. With its tens of billions of unstructured and dynamic pages and its increasing number of users, the World Wide Web poses new great challenges to the existing machine learning algorithms, and at the same time it also fuels the rapid development of new machine learning techniques. This workshop aims at bringing machine learning and web search people together to discuss the fundamental issues in web search from relevance ranking and web spam detection to online advertising.

Topics to be discussed:

- Web page ranking: ranking algorithms and theory, rank aggregation, link analysis
- Online advertising: click-rate prediction, keyword generation, clicks fraud detection, content matching, and auction mechanism
- Usage data: learning from query and click-through logs, interactive experimentation, active learning, exploration/exploitation
- Web spam detection: link spam, content spam, blog spam, cloaking
- Query rewriting: spelling check, query alternation, query suggestion, query classification
- Social networks: online community discovering, trust and reputation, collaborative filtering
- Large-scale machine learning for web search issues

Additional speakers (including submitted talks) will be listed on the website once confirmed.

Please check the website for talk abstracts, updates, and details of short talks.

Machine Learning for Web Search

Friday December 7, 2007

Organizers: Dengyong Zhou, Olivier Chapelle, Thorsten Joachims & Thomas Hofmann

<http://research.microsoft.com/conferences/nips07/websearch.html>

Morning session: 7:30am–10:30am

- 7:30am **TBD**, *Yoram Singer*
- 8:30am **Large Margin Optimization of Ranking Measures**, *Olivier Chapelle, Quoc Le & Alex Smola*
- 8:50am *posters and coffee*
- 9:20am **SoftRank with Gaussian Processes**, *Edward Snelson & John Guiver*
- 9:40am **Manipulation-Resistant Reputations Using Hitting Time**, *John Hopcroft & Daniel Sheldon*
- 10:00am *posters and discussion*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Computational Advertising**, *Andrei Broder*
- 4:30pm **Lightly-Supervised Attribute Extraction for Web Search** *Kedar Bellare, Partha Pratim Talukdar, Giridhar Kumaran, Fernando Pereira, Mark Liberman, Andrew McCallum & Mark Dredze*
- 4:50pm *posters and coffee*
- 5:20pm **Privacy, Personalization, and the Web: A Utility-Theoretic Approach**, *Andreas Krause & Eric Horvitz*
- 5:40pm **Learning to Admit You're Wrong: Statistical Tools for Evaluating Web QA**, *Mark Dredze & Krzysztof Czuba*
- 6:00pm Closing remarks, *the organisers and the invited speakers*

Approximate Bayesian Inference in Continuous/Hybrid Models

Friday December 7, 2007

Matthias Seeger, *Max Plank Institute for Biological Cybernetics*

David Barber, *University College London*

Neil Lawrence, *University of Manchester*

Onno Zoeter, *Microsoft Research*

<http://intranet.cs.man.ac.uk/ai/nips07>

Approximate inference techniques are often fundamental to machine learning successes, and this fast moving field has recently facilitated solutions to large scale problems.

The workshop will provide a forum to discuss unsolved issues, both practical and theoretical, pertaining to the application of approximate Bayesian inference. The emphasis of the workshop will be in characterizing the complexity of inference and the differential strengths and weaknesses of available approximation techniques.

The target audience are practitioners, providing insight into and analysis of problems with certain methods or comparative studies of several methods, as well as theoreticians interested in characterizing the hardness of inference or proving relevant properties of an established method.

Sponsored by the Pascal Network of Excellence and Microsoft Research Cambridge.

Approximate Bayesian Inference in Continuous/Hybrid Models

Friday December 7, 2007

Organizers: Matthias Seeger, David Barber, Neil Lawrence & Onno Zoeter

Posters: throughout the day

Variational Optimisation by Marginal Matching : Lawrence, N.

Eval. of Posterior Approx. Schemes for Gaussian Process Classifiers: Cawley, G., Girolami, M.

Improving on Expectation Propagation : Paquet, U., Winther, O., Opper, M.

Approx. the Part. Function by Deleting and then Correcting for Model Edges : Choi, A., Darwiche, A.

Particle Belief Propagation : Ihler, A.

Morning session: 7:30am–10:30am

7:30am **Introduction**, Matthias Seeger

7:45am **Invited Talk: Using Infer.NET to Compare Inference Algorithms**, John Winn

8:20am *Poster Spotlights*

8:35am **Large-scale Bayesian Inference for Collaborative Filtering**, Ole Winther

8:55am *Coffee Break and Poster Viewing*

9:10am **Invited Talk: TBA**, Manfred Opper

9:45am **A Completed Information Projection Interpretation of Expectation Propagation**,
John Walsh

10:05am *Poster Session*

Afternoon session: 3:30pm–6:30pm

3:30pm **Invited Talk : Approximation and Inference using Latent Variable Sparse Linear Models**, David Wipf

4:05pm **The Probability Hypothesis Density Filter: A Tutorial**, A. Taylan Cemgil

4:35pm **Efficient Bounds for the Softmax Function, Applications to Inference in Hybrid Models**,
Guillaume Bouchard

4:55pm *Coffee Break and Poster Viewing*

5:10pm **Invited Talk: Message-Passing Algorithms for Gaussian Markov Random Fields and Non-Linear Optimization**, Jason K. Johnson

5:45pm **Bounds on the Bethe Free Energy for Gaussian Networks**, Botond Cseke

6:05pm *Wrap-Up. Questions from the Audience to Speakers.*
Poster Viewing (optional)

Beyond Simple Cells

Probabilistic Models for Visual Cortical Processing

Friday December 7, 2007

Richard E. Turner, Pietro Berkes, Maneesh Sahani, *Gatsby Computational Neuroscience Unit*

<http://www.gatsby.ucl.ac.uk/~berkes/docs/NIPS07>

The goal of the workshop is to assess the success of computational models for cortical processing based on probabilistic models and to suggest new directions for future research. Models for simple cells are now well established; how can the field progress beyond them?

We will review important questions in the field through both theoretical and experimental lenses. We have an invited panel of experimentalists to provide the latter with the hope of inspiring new research directions of greater general neuroscientific interest.

More precisely, the issues to be discussed include:

- Moving beyond a two-layer hierarchy: how can we bridge the gap between objects and Gabors?
- What experimental results should we attempt to model? Are current comparisons with physiology at all relevant?
- What aspects of visual input are relevant for modeling (eye movements, head movements, color etc.)? How relevant is time?
- What experimental results would we like to have?
- Is the cortical representation best described by energy models or generative models?
- What inference and learning algorithms should we use?

Modelers will present their current work in short talks. The experimentalists, on the other hand, will form a panel of experts which will drive the discussion, bringing their perspective into play, and constructively expose discrepancies with biology.

Through this process, the aim is to attack the main list of questions above, and determine which aspects of the vision problem would especially benefit from a stronger collaboration between modelers and experimentalists.

Beyond Simple Cells

Probabilistic Models for Visual Cortical Processing

Friday December 7, 2007

Organizers: Richard E. Turner & Pietro Berkes & Maneesh Sahani

A panel composed of three leading experimentalists, Dario Ringach, Jozsef Fiser, and Andreas Tolias, will drive discussion at the end of each talk.

Morning session: 7:30am–10:30am

- 7:30am **Introductory remarks on behalf of the modelers**, *M. Sahani*
- 7:40am **Introductory remarks on behalf of the experimentalists**, *J. Fiser*
- 7:50am **Building and testing multi-stage models for cortical processing**, *E. Simoncelli*
- 8:30am **Natural image statistics and contextual visual processing**, *O. Schwartz*
- 9:10am *coffee break*
- 9:30am **Learning to generalize over regions of natural images**, *M. Lewicki*
- 10:10am *poster session*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Can Markov random fields tell us anything about visual receptive fields?**, *M. Black*
- 4:10pm **Self-taught learning via unsupervised discovery of structure**, *A. Ng*
- 4:50pm *coffee break*
- 5:10pm **What the other 85% of V1 is doing**, *B. Olshausen*
- 5:50pm *discussion*

Workshop on Efficient Machine Learning Overcoming Computational Bottlenecks in Machine Learning

Friday-Saturday December 7-8, 2007

Samy Bengio, *Google*

Corinna Cortes, *Google*

Dennis DeCoste, *Microsoft Live Labs*

Francois Fleuret, *IDIAP Research Institute*

Ramesh Natarajan, *IBM T.J. Watson Research Lab*

Edwin Pednault, *IBM T.J. Watson Research Lab*

Dan Pelleg, *IBM Haifa Research Lab*

Elad Yom-Tov, *IBM Haifa Research Lab*

<http://bigml.wikispaces.com/cfp>

The ever increasing size of available data to be processed by machine learning algorithms has yielded several approaches, from online algorithms to parallel and distributed computing on multi-node clusters. Nevertheless, it is not clear how modern machine learning approaches can either cope with such parallel machineries or take into account strong constraints regarding the available time to handle training and/or test examples. This workshop will explore two alternatives:

- Modern machine learning approaches that can handle real time processing at train and/or at test time, under strict computational constraints (when the flow of incoming data is continuous and needs to be handled)
- Modern machine learning approaches that can take advantage of new commodity hardware such as multicore, GPUs, and fast networks.

This two-day workshop aims to set the agenda for future advancements by fostering a discussion of new ideas and methods and by demonstrating the potential uses of readily-available solutions. It will bring together both researchers and practitioners to offer their views and experience in applying machine learning to large scale learning.

Workshop on Efficient Machine Learning (Part 1)

Friday December 7, 2007

Organizers: Samy Bengio, Corinna Cortes, Dennis DeCoste, Francois Fleuret, Ramesh Natarajan, Edwin Pednault, Dan Pelleg & Elad Yom-Tov

Morning session: 7:30am–10:30am

- 7:30am **Introduction**, *S. Bengio*
- 7:40am **Invited talk: Efficient Learning with Sequence Kernels**, *C. Cortes, Google*
- 8:20am **Fully Distributed EM for Very Large Datasets**, *J. Wolfe, A. Haghighi, D. Klein*
- 8:45am **A Parallel N-Body Data Mining Framework**, *G.F. Boyer, R.N. Riegel, A.G. Gray*
- 9:10am *coffee break*
- 9:25am **Invited talk: Large Scale Clustering and Regression Using the IBM Parallel Machine Learning Toolbox**,
E. Yom-Tov, D. Pelleg, R. Natarajan, E. Pednault, N. Slonim, IBM Research
- 10:05am **Batch Performance for an Online Price**, *K. Crammer, M. Dredze, J. Blitzer, F. Pereira*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Invited talk: Model Compression: Bagging Your Cake and Eating It Too**,
R. Caruana, Cornell, and D. DeCoste, Microsoft
- 4:10pm **Poster Session:**
Fast SVD for Large-Scale Matrices, *M. P. Holmes, A. G. Gray, C. L. Isbell*
Large-Scale Euclidean MST and Hierarchical Clustering, *W. March, A. Gray*
SVM Ocas, *V. Franc and S. Sonnenburg*
A Large-Scale Gaussian Belief Propagation Solver for Kernel Ridge Regression, *D. Bickson, E. YomTov, D. Dolev*
Parallel support vector machine training, *K. Woodsend, J. Gondzio*
- 5:10pm **Invited talk: Architecture Conscious Data Analysis: Progress and Future Outlook**,
S. Parthasarathy, Ohio State University
- 5:50pm **Invited talk: Who is Afraid of Non-Convex Loss Functions?**,
Y. LeCun, Courant Institute, New York University

MLSys: Statistical Learning Techniques for Solving Systems Problems

Friday-Saturday, December 7-8, 2007

Sumit Basu, *Microsoft Research*

Archana Ganapathi, *UC Berkeley*

Emre Kiciman, *Microsoft Research*

Fei Sha, *Yahoo! Research*

<http://radlab.cs.berkeley.edu/MLSys>

In the last few years, there has been a budding interaction between machine learning and computer systems researchers. In particular, statistical machine learning techniques have found a wide range of successful applications in many core systems areas, from designing computer microarchitectures and analyzing network traffic patterns to managing power consumption in data centers and beyond. However, connecting these two areas has its challenges: while systems problems are replete with mountains of data and hidden variables, complex sets of interacting systems, and other exciting properties, labels can be hard to come by, and the measure of success can be hard to define. Furthermore, systems problems often require much more than high classification accuracy - the answers from the algorithms need to be both justifiable and actionable.

Dedicated workshops in systems conferences have emerged (for example, SysML 2006 and SysML 2007) to address this area, though they have had little visibility to the machine learning community. A primary goal of this workshop is thus to expose these new research opportunities in systems areas to machine learning researchers, in the hopes of encouraging deeper and broader synergy between the two communities. During the workshop, through various overviews, invited talks, poster sessions, group discussions, and panels, we would like to achieve three objectives. First, we wish to discuss the unique opportunities and challenges that are inherent to this area. Second, we want to discuss and identify “low-hanging fruit” that are be more easily tackled using existing learning techniques. Finally, we will cover how researchers in both areas can make rapid progress on these problems using existing toolboxes for both machine learning and systems.

We hope that this workshop will present an opportunity for intensive discussion of existing work in machine learning and systems, as well as inspire a new generation of researchers to become involved in this exciting domain.

MLSys: Statistical Learning Techniques for Solving Systems Problems (Part 1)

Friday December 7, 2007

Organizers: Sumit Basu, Archana Ganapathi, Emre Kiciman, & Fei Sha

Morning session: 7:30am–10:30am

- 7:30am **Overview of the Machine Learning and Systems Research Area,**
Sumit Basu, Archana Ganapathi, Emre Kiciman, and Fei Sha
- 8:15 am **Invited Talk: Learning and Systems: The Good, the Bad, and the Ugly...,**
Moises Goldszmidt, Microsoft Research
- 8:45am **Invited Talk: Blind Source Separation in Network Tomography,** *Irina Rish, IBM Research*
- 9:15am *Coffee break*
- 9:30am *Poster Spotlights*
- 10:30am *Adjourn*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Invited Talk: Structure from Failure: Inferring Functional Network Dependencies from Observed Failure Data,** *Thore Graepel, Microsoft Research*
- 4:00pm **Invited Talk: Open Problems in Performance Diagnosis for Distributed File Systems,**
Alice Zheng, Carnegie Mellon University
- 4:30pm *Poster Session*
- 6:30pm *Adjourn*

Posters

- 1 **Reinforcement Learning for Capacity Tuning of Multi-Core Servers**
Liat Ein-Dor¹, Yossi Itach¹, Aharon Bar-Hillel¹, Amir Di-Nur², and Ran Gilad-Bachrach¹
¹Intel Research Labs and ²Intel IT
- 2 **Melody - Expert-Free System Analysis**
Sivan Sabato, Elad Yom-Tov, and Ohad Rodeh
IBM Haifa Research Labs
- 3 **Online Learning with Constraints for Systems**
Branislav Kveton¹, Jia Yuan Yu², Georgios Theodorou¹, and Shie Mannor²
¹Intel Research and ²McGill University
- 4 **Reinforcement Learning for Utility-Based Grid Scheduling**
Julien Perez, Balazs Kegl, and Cecile Germain-Renaud
University of Paris-Sud
- 5 **Learning to Dynamically Allocate Memory**
Nicolo Cesa-Bianchi¹ and Ofer Dekel²
¹Universita delgi Studi di Milano and ²Microsoft Research
- 6 **Software Unit Test Prioritization: A Statistical Approach**
Jason V. Davis and Emmet Witchel
University of Texas at Austin
- 7 **Using Machine Learning to Discover Service and Host Dependencies in Networked Systems**
Paul Barham, Richard Black, Moises Goldszmidt, Rebecca Isaacs, John MacCormick, Richard Mortier, and Aleksandr Simma
Microsoft Research
- 8 **Real-Time Anomaly Detection in Software Appliances**
F. Knorn, D.J. Leith, and R.N. Shorten
Hamilton Institute, Ireland
- 9 **Policy Search Optimization for Spatial Path Planning**
Matthew E. Taylor, Katherine E. Coons, Behnam Robatmili, Doug Burger, and Kathryn S. McKinley
University of Texas at Austin
- 10 **Dynamic Clustering of Interference Domains in Large Scale 802.11 Networks**
Reza Lotun, Kan Cai, and Mike Feeley
University of British Columbia
- 11 **Network Planning Using Reinforcement Learning**
Eric Vigeant, Shie Mannor, and Doina Precup
McGill University
- 12 **Approaches to Anomaly Detection using Host Network-Traffic Traces**
John Mark Agosta, Jaideep Chardrashekar, Frederic Giroire, Carl Livadas, and Jing Xu
Intel Research

- 13 **Approximate Decision Making in Large-Scale Distributed Systems**
Ling Huang¹, Minos Garofalakis², Anthony D. Joseph³, and Nina Taft¹
¹Intel Research, ²Yahoo! Research, and ³UC Berkeley
- 14 **Tracking Malicious Regions of the IP Space**
Avrim Blum, Dawn Song, and Shobha Venkataraman
Carnegie Mellon University
- 15 **Response-Time Modeling for Resource Allocation and Energy-Informed SLAs**
Peter Bodik, Charles Sutton, Armando Fox, David Patterson, and Michael Jordan
UC Berkeley
- 16 **Discovering the Runtime Structure of Software with Probabilistic Generative Models**
Scott Richardson, Michael Otte, Michael C. Mozer, Amer Diwan, Dan Connors
University of Colorado
- 17 **Learning Link Quality in Wireless Communication Networks**
Joseph E. Gonzalez, Andreas Krause, Katherine Chen, and Carlos Guestrin
Carnegie Mellon University

Machine Learning in Computational Biology

Friday-Saturday December 7-8, 2007

Gal Chechik, *Google Research*

Christina Leslie, *Memorial Sloan-Kettering Cancer Center*

William Stafford Noble, *University of Washington*

Gunnar Rätsch, *Friedrich Miescher Laboratory of the Max Planck Society (Tübingen, Germany)*

Quaid Morris, *University of Toronto*

Koji Tsuda, *Max Planck Institute for Biological Cybernetics (Tübingen, Germany)*

<http://www.mlcb.org>

The field of computational biology has seen dramatic growth over the past few years, both in terms of available data, scientific questions and challenges for learning and inference. These new types of scientific and clinical problems require the development of novel supervised and unsupervised learning approaches.

In particular, the field is characterized by a diversity of heterogeneous data. The human genome sequence is accompanied by real-valued gene and protein expression data, functional annotation of genes, genotyping information, a graph of interacting proteins, a set of equations describing the dynamics of a system, localization of proteins in a cell, a phylogenetic tree relating species, natural language text in the form of papers describing experiments, partial models that provide priors, and numerous other data sources.

The goal of this workshop is to present emerging problems and machine learning techniques in computational biology, with a particular emphasis on methods for computational learning from heterogeneous data. The workshop includes invited and submitted talks from experts in the fields of biology, bioinformatics and machine learning. The topics range from case studies of particular biological problems to novel learning approaches in computational biology.

Machine Learning in Computational Biology (Part 1)

Friday December 7, 2007

Gal Chechik, *Google Research*

Christina Leslie, *Memorial Sloan-Kettering Cancer Center*

William Stafford Noble, *University of Washington*

Gunnar Rätsch, *Friedrich Miescher Laboratory of the Max Planck Society (Tübingen, Germany)*

Quaid Morris, *University of Toronto*

Koji Tsuda, *Max Planck Institute for biological Cybernetics (Tübingen, Germany)*

<http://www.mlcb.org>

Morning session – Interactions and function prediction: 7:30am–10:30am

- 7:45am **Gene Function Prediction from Multiple Data Sources Using GO Priors,**
Sara Mostafavi, Debajyoti Ray, David Warde-Farley, Chris Grouios and Quaid Morris
- 8:10am **Kernel Methods for In Silico Chemogenomics,** *Laurent Jacob and Jean-Philippe Vert.*
- 8:35am **Searching for Functional Gene Modules With Interaction component models,**
Juuso Parkkinen and Samuel Kaski
- 9:00am *coffee break*
- 9:15am **Reaction Graph Kernels for Discovering Missing Enzymes in the Plant Secondary Metabolism,** *Hiroto Saigo, Masahiro Hattori and Koji Tsuda*
- 9:40am **An Automated Combination of Kernels for Predicting Protein Subcellular Localization,**
Alexander Zien and Cheng Soon Ong
- 10:05am *Discussion*

Afternoon session – Gene expression and regulation: 3:30pm–6:30pm

- 3:30pm **Growth-specific Programs of Gene Expression,** *Edo Airolidi*
- 3:55pm **Stochastic Multiscale Modeling Methods for Stem Cell Niches,**
Guy Yosiphon, Kimberly Gokoffski, Anne Calof, Arthur Lander and Eric Mjolsness
- 4:20pm *coffee break*
- 4:35pm **Invited talk: Informative Positional Priors for DNA Motif Discovery: Nucleosome Positioning, Sequence Conservation, DNA Stability,** *Alex Hartemink (tentative)*
- 5:20pm **Nucleosome Positioning from Tiling Microarray Data,**
Moran Yassour, Tommy Kaplan, Ariel Jaimovich and Nir Friedman
- 5:45pm **Deconvolution Yields a High-resolution View of Global Gene Expression During the Cell Cycle,**
Allister Bernard, David Orlando, Charles Lin, Edwin Iversen, Steven Haase and Alexander Hartemink
- 6:10pm *Discussion: Future plans for the Computational Biology workshop, The organizers*

Large Scale Brain Dynamics

Friday and Saturday, December 7 and 8, 2007

Kai J. Miller, *University of Washington*

Ryan T. Canolty, *University of California, Berkeley*

http://www.cs.washington.edu/homes/kai/nips_07_workshop.htm

The "Large Scale Brain Dynamics" workshop is focused on the dynamics of electrical activity of the brain on the 5-10 mm scale, with an emphasis on multielectrode electrocorticographic (ECoG) recording. Central questions include: What are the relevant aspects of the large scale cortical signal for feature extraction? Given current clinical technology, what constraints are there on the large scale potential measurements at the brain surface? (That is, is there a cutoff frequency for signal extraction? How reproducible are phenomena? What is the true temporal, spatial, and spectral fidelity?) If two cortical areas are communicating, what sort of ECoG signal features would be present? What is the best way to track variable-delay activity in multiple brain regions associated with a complex task? The primary goal of the workshop is to provide a forum to identify and discuss the key issues of the field.

Large Scale Brain Dynamics (Part 1)

Friday, December 7, 2007

Organizers: Kai J. Miller & Ryan T. Canolty

Morning session: 7:30am–10:30am

- 7:30am **An introduction to the multi-scale nature of brain dynamics**, *Terry Sejnowski*
- 8:00am **Multi-area LFP and flexprint ECoG recordings**, *Pascal Fries*
- 8:35am **ECoG spectral correlates of human behavior**, *Nathan Crone*
- 9:10am *coffee break*
- 9:20am **Cross-frequency coupling in the human hippocampus**, *Florian Mormann*
- 10:27am **The peaks and the power laws in the cortical spectrum**, *Kai Miller*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Relevant features in the LFP**, *Bijan Pesaran*
- 4:05pm **The variety of cell responses in polytrode recordings and their relation to LFP phase**,
Tim Blanche/Kilian Koepsell
- 4:40pm **Optimal information extraction from electrophysiology data**, *Jose Principe*
- 5:15pm *coffee break*
- 5:30pm **Powerpoint poster highlights followed by poster session**,

Brain, Music & Cognition (Part 1): Learning The Structure of Music and its Effects on The Brain

Friday December 7, 2006

David R. Hardoon, *University College London*

Eduardo Reck Miranda, *University of Plymouth*

John Shawe-Taylor, *University College London*

<http://homepage.mac.com/davidrh/MBCworkshop07/>

Until recently, the issue of musical representation had focused primarily on symbolic notation of musical information and structure, and on the representation of musical performance. Research on how we represent musical experience in the brain is emerging as a rich area of investigation thanks to ongoing advances in brain-scanning technology such as EEG and fMRI.

This day of the workshop addresses the problem of representation of musical experience in the brain from a computational modelling approach chiefly based on machine learning and signal processing.

The overarching question addressed by this workshop is whether we can devise efficient methods to study and model the representation of musical experience in the brain by combining traditional forms of musical representation (musical notation, audio, performance, etc.) with brain scanning technology.

This problem is of particular relevance for the successful modelling of cognitive music behaviour, design of interactive music systems, informing techniques in contemporary composition, providing methods to enhance music performance, and even helping music analysis in suggesting ways of listening to music.

Brain, Music & Cognition (Part 1): Learning The Structure of Music and its Effects on The Brain

Friday December 7, 2007

Organizers: David R. Hardoon, Eduardo Reck Miranda and John Shawe-Taylor

Morning session: 7:30am–10:30am

- 7:30am Introduction, *D. R. Hardoon*
- 7:45am **Invited Talk**, *Sebastian Jentschke (TBC)*
- 8:35am **Neural Correlates of Tonality in Music**,
S. Durrant, D. R. Hardoon, E. R. Miranda, J. Shawe-Taylor, A. Brechmann and H. Scheich
- 9:00am *coffee break & poster viewing*
- 9:20am **Music Preference Learning with Partial Information**,
Y. Moh, P. Orbanz and J. M. Buhmann
- 9:30am **Linear Programming Boosting for Classification of Musical Genre**,
T. Diethe and J. Shawe-Taylor
- 9:40am **The Conditionally Independent Voice Model**, *C. Raphael*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Time-Frequency and Synchrony Analysis of Responses to Steady-State Auditory and Musical Stimuli from Multichannel EEG**,
T. M. Rutkowski, J. Dauwels, F. Vialatte, A. Cichocki and D. P. Mandic
- 3:55pm **Finding Musically Meaningful Words by Sparse CCA**,
D. A. Torres, D. Turnbull, L. Barrington, B. K. Sriperumbudur and G. R. G. Lanckriet
- 4:20pm **Can Style be Learned? A Machine Learning Approach Towards ‘Performing’ as Famous Pianists**, *L. Dorard, D. R. Hardoon and J. Shawe-Taylor*
- 4:30pm **Discovering Music Structure via Similarity Fusion**,
J. Arenas-García, E. Parrado-Hernández, A. Meng, L. K. Hansen and J. Larsen
- 4:40pm *coffee break & poster viewing*
- 5:00pm **Information Dynamics and the Perception of Temporal Structure in Music**, *S. Abdallah*
- 5:50pm **A Maximum Likelihood Approach to Multiple Fundamental Frequency Estimation From the Amplitude Spectrum Peaks**, *Z. Duan and C. Zhang*
- 6:00pm **A Generative Model for Rhythms**, *J. Paiemen, Y. Grandvalet, S. Bengio and D. Eck*
- 6:10pm Discussion, *M. Plumbley*

Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives

Friday 7th - Saturday 8th December, 2007

Yael Niv, *Princeton University*

Matthew M. Botvinick, *Princeton University*

Andrew C. Barto, *University of Massachusetts, Amherst*

The aim of this workshop is to discuss current ideas from computer science, psychology and neuroscience regarding learning and control of hierarchically structured behavior. Psychological research has long emphasized that human behavior is hierarchically structured. Indeed, a hierarchical organization of human behavior that matches the hierarchical structure of real-world problems has been the focus of much empirical and theoretical research, and has played a pivotal role in research on organized, goal-directed behavior. Behavioral hierarchy has been of longstanding interest within neuroscience as well, where it has been considered to relate closely to prefrontal cortical function. The prefrontal cortex, which, with its high cognitive functions, remains the most poorly understood area of the brain, has been repeatedly implicated in supporting and executing hierarchical learning and control. In yet a third field, recent developments within machine learning have led to the emergence of ‘hierarchical reinforcement learning’. This line of research has begun investigating in depth how optimal control can learn, and make use of, hierarchical structures, specifically, how hierarchies of skills (also termed options, macros or temporally abstract actions) could be learned and utilized optimally.

This workshop brings together front-line researchers from each of these fields, with the aim of gleaning new insights by integrating knowledge from these somewhat disparate areas of active research. The overarching goal is to facilitate sharing of ideas such as to potentially advance research in one field based on ideas and knowledge from other fields. The coming together of these three communities is especially exciting because, arguably, some of the most profound developments in psychology and neuroscience in the last two decades have stemmed from the use of normative ideas from reinforcement learning in thinking about and studying behavior and the brain. There is thus much promise in forging links between the long legacy of insight into human cognition and the more recent normative study of hierarchical control.

Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives (Part 1)

Friday December 7, 2007

Organizers: Yael Niv, Matthew M. Botvinick & Andrew C. Barto

Morning session: 9:00am–12:00pm

- 9:00am **Welcome and introduction to hierarchical reinforcement learning**, *Y. Niv*
- 9:25am **Hierarchical reinforcement learning and the brain: Potential connections**, *M. Botvinick*
- 9:50am **Questions and Discussion**, *Niv & Botvinick*
- 10:05am *coffee break*
- 10:25am **The hierarchies that underlie routine behavior**, *R. Cooper, N. Ruh & D. Mareschal*
- 10:50am **Cognitive control, hierarchy, and the rostro-caudal organization of the prefrontal cortex**,
D. Badre
- 11:15am **Architecture of central executive functions in the human prefrontal cortex**, *E. Koechlin*
- 11:40am **Panel Discussion**, *Cooper, Badre & Koechlin*

Afternoon session: 3:00pm–6:00pm

- 3:00pm **Learning hierarchical structure in policies**, *B. Marthi, L. Kaelbling & T. Lozano-Perez*
- 3:25pm **Hierarchical assignment of behaviours to subgoals**, *W. Moerman, B. Bakker & M. Wiering*
- 3:50pm **Recognizers: A study in learning how to model temporally extended behaviors**,
J. Frank & D. Precup
- 4:15pm **Panel Discussion**, *Marthi, Bakker & Precup*
- 4:35pm *coffee break*
- 4:55pm **Computational, behavioral and neuro-imaging methods investigating the hierarchical organization of prefrontal cortex and goal-oriented behavior**,
J. Reynolds, T. Braver & R. O'Reilly
- 5:20pm **Flexible shaping: How learning in small steps helps**, *K. Krüger & P. Dayan*
- 5:45pm **Panel Discussion**, *Reynolds & Krüger*

Saturday
December 8, 2007

Workshops on Saturday, December 8th[★]

Machine Learning in Adversarial Environments for Computer Security [9:00-12:00; 2:00–5:00]

Richard Lippmann, Pavel Laskov

Hilton: Sutcliffe A

Statistical Network Models

Kevin Murphy, Lise Getoor, Eric Xing, Raphael Gottardo

Hilton: Cheakamus

The Grammar of Vision: Probabilistic Grammar-based Models for Visual Scene Understanding and Object Categorization

Virginia Savova, Josh Tenenbaum, Leslie Kaelbling, Alan Yuille

Westin: Alpine A–C

Machine Learning and Games: Open Directions in Applying Machine Learning to Games [9:00–12:00; 2:00–5:00]

Joaquín Quiñero Candela, Thore Graepel, Ralf Herbrich

Hilton: Mt Currie N

The Urban Challenge – Perspectives of Autonomous Driving [9:00–12:00; 2:00–5:00]

Sebastian Thrun, Chris Urmson, Raul Rojas

Hilton: Black Tusk

Representations and Inference on Probability Distributions

Kenji Fukumizu, Arthur Gretton, Alex Smola

Hilton: Diamond Head

Mechanisms of Visual Attention [8:30–12:00; 2:00-5:30]

Jillian Fecteau, John Tsotsos, Dirk Walther, Vidhya Navalpakkam

Westin: Glacier

Efficient Machine Learning: Overcoming Computational Bottlenecks in Machine Learning (Part 2)

Samy Bengio, Corinna Cortes, Dennis DeCoste, Francois Fleuret, Ramesh Natarajan, Edwin Pednault, Dan Pelleg, Elad Yom-Tov

Hilton: Mt Currie S

MLSys: Statistical Learning Techniques for Solving Systems Problems (Part 2)

Sumit Basu, Archana Ganapathi, Emre Kiciman, Fei Sha

Hilton: Sutcliffe B

Machine Learning in Computational Biology (Part 2)

Gal Chechik, Christina Leslie, Quaid Morris, William Stafford Noble, Gunnar Rätsch, Koji Tsuda

Westin: Nordic

Large Scale Brain Dynamics (Part 2)

Ryan Canolty, Kai Miller

Westin: Emerald A

Music, Brain and Cognition (Part 2): Models of Sound and Cognition

Hendrik Purwins, Xavier Serra, Klaus Obermayer

Westin: Alpine D–E

Hierarchical Organization of Behavior: Computational, Psychological & Neural Perspectives (Part 2) [9:00–12:00; 3:00–6:00]

Yael Niv, Matthew Botvinick, Andrew Barto

Westin: Callaghan

★ All sessions 7:30–10:30 and 3:30–6:30 unless otherwise indicated.

Machine Learning in Adversarial Environments for Computer Security

Saturday December 8, 2007

Richard Lippmann, *MIT Lincoln Laboratory*

Pavel Laskov, *Fraunhofer Institute FIRST and University of Tuebingen*

<http://mls-nips07.first.fraunhofer.de/>

Computer and network security has become an important research area due to the alarming recent increase in hacker activity motivated by profit and both ideological and national conflicts. Increases in spam, bot-nets, viruses, malware, key loggers, software vulnerabilities, zero-day exploits and other threats contribute to growing concerns about security. In the past few years, many researchers have begun to apply machine learning techniques to these and other security problems. Security, however, is a difficult area because adversaries actively manipulate training data and vary attack techniques to defeat new systems. A main purpose of this workshop is examine adversarial machine learning problems across different security applications to see if there are common problems, effective solutions, and theoretical results to guide future research, and to determine if machine learning can indeed work well in adversarial environments. Another purpose is to initiate a dialog between computer security and machine learning researchers already working on various security applications, and to draw wider attention to computer security problems in the NIPS community.

Machine Learning in Adversarial Environments for Computer Security

Saturday December 8, 2007

Organizers: Richard Lippmann & Pavel Laskov

Morning session: 9:00am–12:00am

9:00am **Opening Remarks.** *R. Lippmann and P. Laskov*

9:10am **Can Machine Learning Be Secure?** *M. Barreno*

9:40am **Foundations of Adversarial Machine Learning.** *D. Lowd, C. Meek, P. Domingos*

10:00am *Poster spotlights*

- **Optimal Spamming: Solving a Family of Adversarial Classification Games.** *M. Brückner, S. Bickel, T. Scheffer*
- **Statistical Classification and Computer Security.** *A.A. Cardenas, J.D. Tygar*
- **Sensor Placement for Outbreak Detection in Computer Security.** *A. Krause, H.B. McMahan, C. Guestrin, A. Gupta*
- **Proactive Vulnerability Assessment of Networks.** *A. Fern, T. Nguyen, S. Dejmal, L.C. Viet*
- **Automatic Detection and Banning of Bontent Stealing Bots for E-commerce.** *N. Poggi, J.L. Berral, T. Moreno, R. Gavalda, J. Torres*
- **Online Training and Sanitization of AD Systems.** *G.F. Cretu, A. Stavrou, M.F. Locasto, S.J. Stolfo*
- **Combining Multiple One-class Classifiers for Hardening Payload-based Anomaly detection Systems.** *R. Perdisci, G. Gu, W. Lee*
- **Using the Dempster-Shafer Theory for Network Traffic Labelling.** *F. Gargiulo, C. Mazzariello, C. Sansone*

10:15pm *Coffee break and poster preview*

10:45am **Content-based Anomaly Detection in Intrusion Detection.** *S.J. Stolfo*

11:15am **A “Poisoning” Attack Against Online Anomaly Detection.** *M. Kloft, P. Laskov*

11:30am *Discussion:*

- When is adversarial machine learning needed?
- When is it effective, are there some convincing examples?
- How does theory help develop better learning methods?

Afternoon session: 2:00pm–5:00pm

2:00pm **Spam, Phishing, Scam: How to Thrive on Fraud and Deception.** *T. Scheffer*

2:30pm **The War Against Spam: A Report From the Front Line.** *B. Taylor, D. Fingal, D. Aberdeen*

2:42pm **Machine Learning-Assisted Binary Code Analysis.**
N. Rosenblum, X. Zhu, B. Miller, K. Hunt

3:00pm *Poster spotlights*

- **Learning to Predict Bad Behavior.** *N. Syed, N. Feamster, A. Gray*
- **Attacking SpamBayes: Compromising a Statistical Spam Filter.** *M. Barreno, F. Chi, A. Joseph, B. Nelson, B. Rubinstein, U. Saini, C. Sutton, D. Tygar, K. Xia*
- **Using visual and semantic features for anti-spam filters.** *F. Gargiulo, A. Penta, A. Picariello, C. Sansone*
- **Supervised Clustering for Spam Detection in Data Streams.** *U. Brefeld, P. Haider, T. Scheffer*
- **Image Spam Filtering by Detection of Adversarial Obfuscated Text.** *F. Roli, B. Biggio, G. Fumera, I. Pillai, R. Satta*
- **Lightweight Hierarchical Network Traffic Clustering.** *A. Hijazi, H. Inoue, A. Somayaji*
- **Intrusion Detection in Computer Systems as a Pattern Recognition Task in Adversarial Environment: a Critical Review.** *I. Corona, G. Giacinto, F. Roli*
- **Learning from a Flaw in a Naive-Bayes Masquerade Detector.** *K. Killourhy, R. Maxion*

3:15pm *Coffee break and poster preview*

3:45pm **Misleading Automated Worm Signature Generators.** *W. Lee*

4:15pm **ALADIN: Active Learning for Statistical Intrusion Detection.**
J. Stokes, J. Platt, J. Kravis, M. Shilman

4:30pm *Discussion:*

- What do we learn about adversarial learning that can be applied across many applications?
- What types of applications can benefit the most?
- What new theory is necessary?

Statistical Models of Networks

Saturday December 8, 2007

Kevin Murphy, *University of British Columbia*

Raphael Gottardo, *University of British Columbia*

Eric Xing, *Carnegie Mellon University*

Lise Getoor, *University of Maryland*

<http://www.cs.ubc.ca/~murphyk/nips07NetworkWorkshop/index.html>

The purpose of the workshop is to bring together people from different disciplines - computer science, statistics, biology, physics, social science, etc - to discuss foundational issues in the modeling of network and relational data. In particular, we hope to discuss various open research issues, such as

- How to represent graphs at varying levels of abstraction, whose topology is potentially condition-specific and time-varying
- How to combine techniques from the graphical model structure learning community with techniques from the statistical network modeling community
- How to integrate relational data with other kinds of data (e.g., gene expression, sequence or text data)

Statistical Models of Networks

Saturday December 8, 2007

Organizers: Kevin Murphy, Raphael Gottardo, Eric Xing and Lise Getoor

Morning session: 7:30am–10:30am

- 7.30am Opening remarks
- 7.35am **Invited talk 1: Statistical Challenges in Network Modelling**, *Stephen Fienberg*
- 8.15am **Invited talk 2**, *Mark Handcock*
- 8.55am Coffee break
- 9.15am Poster spotlights 1
- 9.35am **Invited talk 3: Hierarchical Eigenmodels for Pooling Relational Data**, *Peter Hoff*
- 10.15am Poster spotlights 2

Afternoon session: 3:30pm–6:30pm

- 3.30pm **Invited talk 4: Relational Latent Class Models**, *Volker Tresp*
- 4.10pm **Invited talk 5: Functions and Phenotypes by Integrative Network Analysis**, *Jasmine Zhou*
- 4.50pm Coffee break
- 5.10pm **Invited talk 6: Recent (and Past) Statistical Models for Network Science**,
Stanley Wasserman
- 5.50pm Discussion

The Grammar of Vision: Probabilistic Grammar-Based Models for Visual Scene Understanding and Object Categorization

Saturday December 8, 2007

Virginia Savova, *MIT*

Josh Tenenbaum, *MIT*

Leslie Kaelbling, *MIT*

Alan Yuille, *UCLA*

<http://web.mit.edu/cocosci/nips07.html>

The human ability to acquire a visual concept from a few examples, and to recognize instances of that concept in the context of a complex scene, poses a central challenge to the fields of computer vision, cognitive science, and machine learning. Representing visual objects and scenes as the human mind does is likely to require structural sophistication, something akin to a grammar for image parsing, with multiple levels of hierarchy and abstraction, rather than the "flat" feature vectors which are standard in most statistical pattern recognition. Grammar-based approaches to vision have been slow to develop, largely due to the absence of effective methods for learning and inference under uncertainty. However, recent advances in machine learning and statistical models for natural language have inspired a renewed interest in structural representations of visual objects, categories, and scenes. The result is a new and emerging body of research in computational visual cognition that combines sophisticated probabilistic methods for learning and inference with classical grammar-based approaches to representation. The goal of our workshop is to explore these new directions, in the context of several interdisciplinary connections that converge distinctively at NIPS. We will focus on these challenges: How can we learn better probabilistic grammars for machine vision by drawing on state-of-the-art methods in statistical machine learning or natural language learning? What can probabilistic grammars for machine vision tell us about human visual cognition? How can human visual cognition inspire new developments in computational vision and machine learning?

The Grammar of Vision: Probabilistic Grammar-Based Models for Visual Scene Understanding and Object Categorization

Saturday December 8, 2007

Organizers: Virginia Savova, Josh Tenenbaum, Leslie Kaelbling & Alan Yuille

Morning session: 7:30am–10:30am

- 7:30am **Introduction**, *Josh Tenenbaum*
- 7:40am **Compositionality in Vision**, *Stuart Geman*
- 8:10am **Stochastic Spatio-Temporal Grammars for Images and Video**, *Jeff Siskind*
- 8:40am **Models and Algorithms for Image Parsing**, *Pedro Felszenswalb and David McAllester*
- 9:10am *coffee break*
- 9:30am **Object Categorization: Modeling, Learning and Parsing with Stochastic Image Grammars**, *Song-Chun Zhu*
- 10:00am **Pros and Cons of Grammar-based Approaches to Vision**, *Chris Williams*
- 10:10am **Building 3-D Models from a Single Still Image**, *Ashutosh Saxena and Andrew Ng*
- 10:20am **Video Deconstruction: Revealing Narrative Structure Through Image and Text Alignment**, *Timothee Cour and Ben Taskar*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Unsupervised Structure Learning of a Hierarchical Object Model**,
Leo (Long) Zhu and Alan Yuille
- 4:00pm **Learning Grammatical Models for 2D Object Recognition**,
Margaret Aycinena, Leslie Kaelbling, Tomas Lozano-Perez
- 4:30pm **Is a Sentence Like a House? Generalization Over Visual Categories with Repetitive and Recursive Structure**, *Virginia Savova*
- 5:00pm *coffee break*
- 5:20pm **Acquiring Complex Structured Representations of Object and Scenes by Humans**,
Jozsef Fiser
- 5:50pm *Panel Discussion: Chris Williams, Bruno Olshausen, Antonio Torralba, Josh Tenenbaum*

Machine Learning And Games: MALAGA

Saturday December 8, 2007

Joaquin Quiñonero Candela, *Applied Games, Microsoft Research*

Thore Graepel, *Applied Games, Microsoft Research*

Ralf Herbrich, *Applied Games, Microsoft Research*

<http://research.microsoft.com/mlp/apg/malaga.aspx>

Computer games sales are three time larger than industry software sales, and on par with Hollywood box office sales. Modern computer games are often based on extremely complex simulations of the real world and constitute one of the very few real fields of application for artificial intelligence encountered in everyday live.

Surprisingly, machine learning methods are not present in the vast majority of computer games. There have been a few recent and notable successes in turn-based two-player, discrete action space games such as Backgammon, Checkers, Chess and Poker. However, these successes are in stark contrast to the difficulties still encountered in the majority of computer games, which typically involve more than two agents choosing from a continuum of actions in complex artificial environments. Typical game AI is still largely built around fixed systems of rules that often result in implausible or predictable behaviour and poor user experience.

The purpose of this workshop is to involve the NIPS community in the exciting challenges that games - ranging from traditional table top games to cutting-edge console and PC games - offer to machine learning.

Machine Learning And Games: MALAGA

Saturday December 8, 2007

Organizers: Joaquin Quiñonero Candela, Thore Graepel & Ralf Herbrich

Morning session: 9:00am–12:00am

- 9:00am **Opening**, *Organizers*
- 9:15am **Artificial Intelligence in Halo 3**, *Max Dyckhoff*, Bungie Studios
- 9:45am **Computer (and Human) Perfection at Checkers**, *Jonathan Schaeffer*, University of Alberta
- 10:15am *coffee break*
- 10:30am **State-of-the-art in Computer Poker**, *Michael Bowling*, University of Alberta
- 11:00am **Monte-Carlo Planning for the Game of Go**,
Olivier Teytaud, TAO(Inria, Univ. Paris-Sud, UMR CNRS-8623)
- 11:30am Demos

Afternoon session: 2:00pm–5:00pm

- 2:00pm **Planning in Real-Time Strategy Games**, *Michael Buro*, University of Alberta
- 2:30pm **Learning Driving Tasks in TORCS Using Reinforcement Learning**,
A. Lazaric, D. Loiacono, A. Prete, M. Restelli, P.L. Lanzi, Politecnico di Milano
- 2:50pm **Modelling Go Positions with Planar CRFs**,
Dmitry Kamenetsky, Nicol N. Schraudolph, Simon Günter, S.V.N. Vishwanathan, NICTA
- 3:10pm **Opportunities for Machine Learning to Impact Interactive Narrative**,
David L. Roberts, Mark Riedl, Charles L. Isbell, Georgia Institute of Technology
- 3:30pm *coffee break*
- 3:45pm **Game Development with XNA**, *Chris Satchell*, Microsoft

The Urban Challenge - Perspectives of Autonomous Driving

Saturday December 8, 2007

Sebastian Thrun, *Stanford University*

Chris Urmson, *Carnegie Mellon University*

Raul Rojas, *Freie Universitaet Berlin*

Will Uther, *NICTA*

The invited speakers for this workshop will analyze the results of the Urban Challenge race of robotic cars held on 3rd November in Victorville, California. They will provide details of their teams' software and hardware architectures and will also discuss future developments in the field of autonomous vehicles. Videos and animations of the actual race will be shown. Representatives of the six teams which finished the race, as well as from other teams, will be present.

The exact program and schedule will be posted at the NIPS conference. Here we list only the confirmed speakers and their team's name.

The Urban Challenge - Perspectives of Autonomous Driving

Saturday December 8, 2007

Organizers: Sebastian Thrun, Chris Urmson, Raul Rojas & Will Uther

Morning session: 9:00–12:00

Stanford Racing Team, *M. Montemerlo*

Team Annieway – University of the Armed Forces Munich, *J. Wuensche*

Team Lux - Hamburg, *H. Salow*

Team Cornell, *S. Lupashin, D. Huttenlocher, M. Campbell*

UT Austin, *P. Stone, P. Beeson*

Afternoon session: 2:00–5:00

Tartan Racing - Carnegie Mellon University, *C. Urmson*

Team MIT, *Albert Huang*

Victor Tango - Virginia Tech, *TBA*

Stanford Racing Team - Toyota Technical Center, *Dmitri Dolgov*

Berkeley-Sydney Team, *Will Uther*

Closing remarks, *Organizers*

Representations and Inference on Probability Distributions

Saturday December 8, 2006

Kenji Fukumizu, *Institute for Statistical Mathematics*

Arthur Gretton, *MPI for Biological Cybernetics*

Alexander Smola, *NICTA, Australian National University*

<http://nips2007.kyb.tuebingen.mpg.de/>

When dealing with distributions it is in general infeasible to estimate them explicitly in high dimensional settings, since the associated learning rates can be arbitrarily slow. On the other hand, a great variety of applications in machine learning and computer science require distribution estimation and/or comparison. Examples include testing for homogeneity (the "two-sample problem"), independence, and conditional independence, where the last two can be used to infer causality; data set squashing / data sketching / data anonymisation; domain adaptation (the transfer of knowledge learned on one domain to solving problems on another, related domain) and the related problem of covariate shift; message passing in graphical models (EP and associated algorithms); compressed sensing; and links between divergence measures and loss functions.

The purpose of this workshop is to bring together statisticians, machine learning researchers, and computer scientists working on representations of distributions for various inference and testing problems, to discuss the compromises necessary in obtaining useful results from finite data. In particular, what are the capabilities and weaknesses of different distribution estimates and comparison strategies, and what negative results apply?

Contributors:

- Yasemin Altun, MPI for Biological Cybernetics
- Tugkan Batu, LSE
- Shai Ben-David, University of Waterloo
- Gerard Biau, Paris VI University
- Carlos Guestrin, CMU
- Zaid Harchaoui, ENST
- Piotr Indyk, MIT
- John Shawe-Taylor, UCL

Additional speakers (including submitted talks) will be listed on the website once confirmed.

Please check the website for talk abstracts, updates, and details of short talks.

Representations and Inference on Probability Distributions

Saturday December 8, 2007

Organizers: Kenji Fukumizu, Arthur Gretton & Alex Smola

Morning session: 7:30am–10:30am

- 7:30am **Property Testing of Distributions**, *Tugkan Batu*
- 8:15am **On a L1-test Statistic of Homogeneity**, *Gérard Biau*
- 9:00am *posters and coffee*
- 9:20am **Efficient, Adaptive Inference for Distributions on Permutations**, *Carlos Guestrin*
- 9:45am **Short talks**, *see website*
- 10:20am *posters and discussion*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **A Subjective Measure for Distribution Similarity**, *Shai Ben-David*
- 4:15pm **Sketching and Streaming for Distributions**, *Piotr Indyk*
- 5:00pm *posters and coffee*
- 5:20pm **Short talks**, *see website*
- 6:20pm *Closing remarks, the organisers*

Models and Mechanisms of Visual Attention: A Critical Appraisal

Saturday December 8, 2007

J. Fecteau, *Netherlands Institute of Neurosciences*

D. B. Walther, *Beckman Institute, University of Illinois at Urbana-Champaign*

V. Navalpakkam, *Division of Biology, California Institute of Technology*

J. Tsotsos, *York University, Toronto, Canada*

Visual attention is a basic cognitive process – it reflects the ability to select and process one object or small region of space in the visual scene. Albeit basic, the concept of visual attention has garnered much interest and many different computational theories have been proposed to explain its neural origins. In this workshop, we will celebrate the strengths and examine the weaknesses of three different theories of attention. Following a brief tutorial lecture, we will debate the relative merits the Biased Competition Model of Attention, the Saliency Map of Attention, and the Selective Tuning Model of Attention. Through this debating process, we will consider the strengths and weaknesses of alternative models as well. Finally, we will debate the computational principles that underlie these models, such as the relevance of information theory, signal detection theory, and optimality principles in capturing the influence of visual attention in neural activity and behavior. Indeed, how might dynamical systems be designed to model attentive behavior? The unorthodox organization of this workshop will encourage candid discussions – the invited speakers will discuss the merits of one of these models and/or the weaknesses of another. We hope you join us in this venture.

Models and Mechanisms of Visual Attention: A Critical Appraisal

Saturday December 8, 2007

Organizers: J. Fecteau, D. B. Walther, V. Navalpakkam & J. Tsotsos

Morning session: 8:30am–12:10pm

8:30am **Overview of Attention and Modeling Research,**
Jillian Fecteau, John Tsotsos (breakfast served)

9:00am Biased Competition

9:00am **Neuropsychological evidence of biased competition,** *Shaun Vecera*
9:20am **Biased competition: Where is the bias coming from?,** *Andrew Rossi*
9:40am **Limitations and inaccuracies of biased competition,**
Open floor discussion, led by Julio Martinez-Trujillo

10:00am Break

10:10am Saliency Map

10:10am **Computational principles of the saliency map,** *Laurent Itti*
10:30am **Brain and behavioral evidence of saliency map,** *Kirk Thompson*
10:50am **Limitations and inaccuracies of saliency map,**
Open floor discussion, led by (to be announced)

11:10am Selective Tuning Model

11:10pm **Computational principles of the selective tuning model,** *John Tsotsos*
11:30pm **Brain and behavioral evidence of selective tuning model,** *Max Hopf*
11:50pm **Limitations and inaccuracies of selective tuning model,**
Open floor discussion, led by (to be announced)

Models and Mechanisms of Visual Attention: A Critical Appraisal

Saturday December 8, 2007

Organizers: J. Fecteau, D. B. Walther, V. Navalpakkam & J. Tsotsos

Afternoon session: 2:00pm–5:30pm

Computational Principles for Neurobiological Modeling of Attention

- 2:00pm **Complexity theory: A critical computational principle in modeling attention,** *John Tsotsos*
- 2:20pm **Signal detection theory: The utility of the ideal observer in modeling attention,**
Bill Geisler
- 2:40pm **Information theory: Using Shannon’s information and Bayesian surprise to model
attention,** *Laurent Itti*
- 3:00pm Break
- 3:10pm **Optimal design: The use of optimality principles in modeling attention,**
Vidhya Navalpakkam
- 3:30pm **Role of context: The use of Bayesian theory to model influence of context in attention,**
Antonio Torralba
- 3:50pm **Deploying attention: The use of focal attention to serialize the recognition of multiple
objects,** *Dirk B. Walther*
- 4:10pm **TBA,** *Søren Kyllingsbæk*
- 4:30pm Break
- 4:45pm Open Discussion and Closing Remarks

Workshop on Efficient Machine Learning Overcoming Computational Bottlenecks in Machine Learning

Friday-Saturday December 7-8, 2007

Samy Bengio, *Google*

Corinna Cortes, *Google*

Dennis DeCoste, *Microsoft Live Labs*

Francois Fleuret, *IDIAP Research Institute*

Ramesh Natarajan, *IBM T.J. Watson Research Lab*

Edwin Pednault, *IBM T.J. Watson Research Lab*

Dan Pelleg, *IBM Haifa Research Lab*

Elad Yom-Tov, *IBM Haifa Research Lab*

<http://bigml.wikispaces.com/cfp>

The ever increasing size of available data to be processed by machine learning algorithms has yielded several approaches, from online algorithms to parallel and distributed computing on multi-node clusters. Nevertheless, it is not clear how modern machine learning approaches can either cope with such parallel machineries or take into account strong constraints regarding the available time to handle training and/or test examples. This workshop will explore two alternatives:

- Modern machine learning approaches that can handle real time processing at train and/or at test time, under strict computational constraints (when the flow of incoming data is continuous and needs to be handled)
- Modern machine learning approaches that can take advantage of new commodity hardware such as multicore, GPUs, and fast networks.

This two-day workshop aims to set the agenda for future advancements by fostering a discussion of new ideas and methods and by demonstrating the potential uses of readily-available solutions. It will bring together both researchers and practitioners to offer their views and experience in applying machine learning to large scale learning.

Workshop on Efficient Machine Learning (Part 2)

Saturday December 8, 2007

Organizers: Samy Bengio, Corinna Cortes, Dennis DeCoste, Francois Fleuret, Ramesh Natarajan, Edwin Pednault, Dan Pelleg & Elad Yom-Tov

Morning session: 7:30am–10:30am

- 7:30am **Learning with Millions of Examples and Dimensions - Competition proposal,**
 S. Sonnenburg, Fraunhofer FIRST
- 7:55am **Large Scale Learning with String Kernels,** *S. Sonnenburg, K. Rieck, G. Ratsch*
- 8:20am **Invited talk: Speeding up stochastic gradient descent,** *Y. Bengio, University of Montreal*
- 9:00am *coffee break*
- 9:25am **Invited talk: Stationary Features and Folded Hierarchies for Efficient Object Detection,**
 D. Geman, John Hopkins University
- 10:05am **Efficient Machine Learning Using Random Projections,**
 C. Hegde, M.A. Davenport, M.B. Wakin, R.G. Baraniuk

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Invited talk: New Quasi-Newton Methods for Efficient Large-Scale Machine Learning,**
 S.V.N. Vishwanathan and N. N. Schraudolph, National ICT Australia
- 4:10pm **Online Classification of Large Numerical Data Sets,** *P. Haffner*
- 4:35pm **Poster session:**
 Efficient Off-line and On-line Algorithms Training Conditional Random Fields Approximating
 Jacobians, C.-N. Hsu, H.-S. Huang, Y.-M. Chang
 Online Stacked Graphical Learning, Z. Kou, V. R. Carvalho, W. W. Cohen
 High Detection-rate Cascades for Real-Time Object Detection, H. Masnadi-Shirazi, N.
 Vasconcelos
 A Stochastic Quasi-Newton Algorithm for Non-Smooth and Non-Convex Optimization, S.
 Gunter, J. Yu, P. Sunehag, N. N. Schraudolph
- 6:00pm **Large Scale Sequence Labelling,** *A. Bordes, L. Bottou*
- 6:25pm **Wrap-up**

MLSys: Statistical Learning Techniques for Solving Systems Problems

Friday-Saturday, December 7-8, 2007

Sumit Basu, *Microsoft Research*

Archana Ganapathi, *UC Berkeley*

Emre Kiciman, *Microsoft Research*

Fei Sha, *Yahoo! Research*

<http://radlab.cs.berkeley.edu/MLSys>

In the last few years, there has been a budding interaction between machine learning and computer systems researchers. In particular, statistical machine learning techniques have found a wide range of successful applications in many core systems areas, from designing computer microarchitectures and analyzing network traffic patterns to managing power consumption in data centers and beyond. However, connecting these two areas has its challenges: while systems problems are replete with mountains of data and hidden variables, complex sets of interacting systems, and other exciting properties, labels can be hard to come by, and the measure of success can be hard to define. Furthermore, systems problems often require much more than high classification accuracy - the answers from the algorithms need to be both justifiable and actionable.

Dedicated workshops in systems conferences have emerged (for example, SysML 2006 and SysML 2007) to address this area, though they have had little visibility to the machine learning community. A primary goal of this workshop is thus to expose these new research opportunities in systems areas to machine learning researchers, in the hopes of encouraging deeper and broader synergy between the two communities. During the workshop, through various overviews, invited talks, poster sessions, group discussions, and panels, we would like to achieve three objectives. First, we wish to discuss the unique opportunities and challenges that are inherent to this area. Second, we want to discuss and identify “low-hanging fruit” that are be more easily tackled using existing learning techniques. Finally, we will cover how researchers in both areas can make rapid progress on these problems using existing toolboxes for both machine learning and systems.

We hope that this workshop will present an opportunity for intensive discussion of existing work in machine learning and systems, as well as inspire a new generation of researchers to become involved in this exciting domain.

MLSys: Statistical Learning Techniques for Solving Systems Problems (Part 2)

Saturday December 8, 2007

Organizers: Sumit Basu, Archana Ganapathi, Emre Kiciman, & Fei Sha

Morning session: 7:30am–10:30am

- 7:30am **Panel: Building Blocks, Tools, and Datasets for MLSys Problems,**
 Workshop Organizers and Guests
- 8:30am **Brainstorming: Short and Long Term Visions for MLSys,**
 Workshop Participants (in small groups)
- 9:30am *Coffee Break*
- 9:45am **Presentation of Brainstorming Results,** *Workshop Participants*
- 10:30am *Adjourn*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Invited Talk: Inferring Bugs in Systems Code with Graphical Models and Static**
 Program Analysis, *Ted Kremenek, Stanford University*
- 4:00pm **Invited Talk: Beyond 'grep': Learning from Event Logs,** *Ira Cohen, HP Labs*
- 4:30pm *Coffee Break*
- 4:45pm **Panel: Open Problems and Challenges in MLSys,** *Workshop Organizers and Guests*
- 5:45pm **Wrap-Up and Compilation of an MLSys Reading List,** *Workshop Participants*
- 6:30pm *Adjourn*

Machine Learning in Computational Biology

Friday-Saturday December 7-8, 2007

Gal Chechik, *Google Research*

Christina Leslie, *Memorial Sloan-Kettering Cancer Center*

William Stafford Noble, *University of Washington*

Gunnar Rätsch, *Friedrich Miescher Laboratory of the Max Planck Society (Tübingen, Germany)*

Quaid Morris, *University of Toronto*

Koji Tsuda, *Max Planck Institute for Biological Cybernetics (Tübingen, Germany)*

<http://www.mlcb.org>

The field of computational biology has seen dramatic growth over the past few years, both in terms of available data, scientific questions and challenges for learning and inference. These new types of scientific and clinical problems require the development of novel supervised and unsupervised learning approaches.

In particular, the field is characterized by a diversity of heterogeneous data. The human genome sequence is accompanied by real-valued gene and protein expression data, functional annotation of genes, genotyping information, a graph of interacting proteins, a set of equations describing the dynamics of a system, localization of proteins in a cell, a phylogenetic tree relating species, natural language text in the form of papers describing experiments, partial models that provide priors, and numerous other data sources.

The goal of this workshop is to present emerging problems and machine learning techniques in computational biology, with a particular emphasis on methods for computational learning from heterogeneous data. The workshop includes invited and submitted talks from experts in the fields of biology, bioinformatics and machine learning. The topics range from case studies of particular biological problems to novel learning approaches in computational biology.

Machine Learning in Computational Biology (Part 2)

Saturday December 8, 2007

Gal Chechik, *Google Research*

Christina Leslie, *Memorial Sloan-Kettering Cancer Center*

William Stafford Noble, *University of Washington*

Gunnar Rätsch, *Friedrich Miescher Laboratory of the Max Planck Society (Tübingen, Germany)*

Quaid Morris, *University of Toronto*

Koji Tsuda, *Max Planck Institute for biological Cybernetics (Tübingen, Germany)*

<http://www.mlcb.org>

Morning session – Proteomics and Imaging: 7:30am–10:30am

- 7:45am **Markov Random Field Based Automatic Image Alignment for Electron Tomography,**
Farshid Moussavi, Fernando Amat, Mark Horowitz, Gal Elidan, Luis Comolli and Kenneth Downing
- 8:10am **Spatial Querying of Imaging Mass Spectrometry Data: A Nonnegative Least Squares Approach,**
Raf Van de Plas, Kristiaan Pelckmans, Bart De Moor and Etienne Waelkens
- 8:35am **Automatic mining of whole cell tomograms for cancer detection,**
Rajesh Narasimha, Hua Ouyang, Alexander Gray, Steven W. McLaughlin and Sriram Subramaniam
- 9:00am *coffee break*
- 9:15am **A Generalized Framework for Protein Sequence Annotation,**
Huzefa Rangwala, Chris Kauffman and George Karypis
- 9:40am **Statistical learning of peptide retention behavior in chromatographic separations: A new kernel-based approach for computational proteomics,**
Nico Pfeifer, Andreas Leinenbach, Christian Huber and Oliver Kohlbacher
- 10:05am *Discussion*

Afternoon session – Genetic variation and evolution: 3:30pm–6:30pm

- 3:30pm **A phylogenetically-corrected dependency network for HIV immune-adaptation,**
Jonathan Carlson, C. Kadie, Z. Brumme, C. Brumme, P. R. Harrigan and D. Heckerman
- 3:55pm **Machine Learning Algorithms for Polymorphism Detection,**
George Zeller, G. Schweikert, R. Clark, S. Ossowski, P. Shin, . . . , B. Schölkopf and G. Rätsch
- 4:20pm *coffee break*
- 4:35pm **Invited talk: Natural history and evolutionary principles of gene duplication in fungi,**
Nir Friedman (tentative)
- 5:20pm **A finite state transducer approach to haplotype phasing,** *Sean O’Rourke and Eleazar Eskin*
- 5:45pm **Positional Oligomer Importance Matrices,**
Soeren Sonnenburg, Alexander Zien, Petra Philips and Gunnar Rätsch
- 6:10pm *Closing remarks, The organizers*

Large Scale Brain Dynamics

Friday and Saturday, December 7 and 8, 2007

Kai J. Miller, *University of Washington*

Ryan T. Canolty, *University of California, Berkeley*

http://www.cs.washington.edu/homes/kai/nips_07_workshop.htm

The "Large Scale Brain Dynamics" workshop is focused on the dynamics of electrical activity of the brain on the 5-10 mm scale, with an emphasis on multielectrode electrocorticographic (ECoG) recording. Central questions include: What are the relevant aspects of the large scale cortical signal for feature extraction? Given current clinical technology, what constraints are there on the large scale potential measurements at the brain surface? (That is, is there a cutoff frequency for signal extraction? How reproducible are phenomena? What is the true temporal, spatial, and spectral fidelity?) If two cortical areas are communicating, what sort of ECoG signal features would be present? What is the best way to track variable-delay activity in multiple brain regions associated with a complex task? The primary goal of the workshop is to provide a forum to identify and discuss the key issues of the field.

Large scale brain dynamics (Part 2)

Saturday, December 8, 2007

Organizers: Kai J. Miller & Ryan T. Canolty

Morning session: 7:30am–10:30am

- 7:30am **Clinic issues and functional mapping with ECoG**, *Jeffrey Ojemann*
- 8:05am **Feedback based changes in the cortical distribution of ECoG signals**, *Gerwin Schalk*
- 8:40am **From theory to the clinic: Translational research and clinical considerations for field potentials**, *Nathan Crone*
- 9:15am *coffee break*
- 9:25am **Spatiotemporal dynamics of word processing in the human cortex**, *Ryan Canolty*
- 9:55am **Scale-free neocortical dynamics**, *Walter Freeman*

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Pattern classification in brain signals**, *Klaus Mueller*
- 4:05pm **Comparing signals: electroencephalography, electromeningiography, and electrocorticography**, *Brad Voytek*
- 4:40pm **Brain oscillations control timing of single-neuron activity in humans**, *Joshua Jacobs*
- 5:15pm *coffee break*
- 5:25pm **Internally generated assembly sequences in the hippocampus and episodic memory**, *Eva Pastalkova*
- 6:00pm **Spatiotemporal Brain Dynamics**, *Scott Makeig*
- 6:30pm **General Discussion and Summary**,

Music, Brain & Cognition (Part 2): Models of Sound and Music Cognition

Saturday December 8, 2007

Hendrik Purwins, *Universitat Pompeu Fabra*

Xavier Serra, *Universitat Pompeu Fabra*

Klaus Obermayer, *Technische Universität Berlin*

<http://homepage.mac.com/davidrh/MBCworkshop07/Workshop.html>

<http://iua-share.upf.es/wikis/mbc2007>

During the second day, the workshop will focus on the modeling of sound, music perception and cognition. These can provide, with the crucial role of machine learning, a break-through in various areas of music technology, in particular: Music Information Retrieval (MIR), expressive music synthesis, interactive music making, and sound design. Understanding of music cognition in its implied top-down processes can help to decide which of the many descriptors in MIR are crucial for the musical experience and which are irrelevant.

In particular music cognition models of three phenomena are of interest: 1) Emergence of categories in music and sound perception, 2) Attention as a top-down mechanism of modulating perception, 3) Goal-directed action in sound and music generation. How can machine learning be used for simulating these processes? Graphical models, recurrent networks, clustering and visualization methods have been suggested so far. In some instances the question arises how modeling approaches from language and vision can be transferred/adapted to music.

Categorization, attention, and goal-directedness are interrelated. The goal of an action determines the attentional selectivity. To a great extent, categorization is based on the interaction with the object or the instrument. We expect that cognitively informed machine learning methods or integrated cognition models can help to solve crucial problems in music cognition and music technology.

The target group is of researchers within the fields of (Music) Cognition, Music Technology, Machine Learning, Psychology, Sound Design, Signal Processing and Brain Imaging.

The organisers of the workshop are investigators for three main European projects; Learning the Structure of Music (Le StruM), Closing the Loop of Evaluation and Design (CLOSED), Emergent Cognition through Active Perception (EmCAP).

Music, Brain & Cognition (Part 2): Models of Sound and Music Cognition

Saturday December 8, 2007

Organizers: Hendrik Purwins, Xavier Serra & Klaus Obermayer

Morning session: 7:30am–10:30am

- 7:30am **Title to be announced, D. Huron**
- 8:10am **The Importance of Cognitive and Motor Constraints in Modeling Music, D. Eck**
- 8:50am *coffee break and poster viewing:*
Adaptive Control for Sound Design, K. Adiloglu, Y. Visell
Comparison of Biologically Inspired Classification Schemes for Everyday Sounds,
R. Annie, E. Martinez, K. Adiloglu, H. Purwins, K. Obermayer
An Assessment of Sound Annoyance as a Function of Consonance, S. H. Chon
Singing Voices Synthesis and Gestural Control for Mimophony, Gualtieri
- 9:10am **Title to be announced, N.N.**
- 9:50am **Hierarchical Bayesian Models for Audio and Music Processing, A. T. Cemgil**

Afternoon session: 3:30pm–6:30pm

- 3:30pm **Title to be announced, M. Slaney**
- 4:10pm **An Auditory Model for the Detection of Perceptual Onsets and Beat Tracking in Singing,**
M. Coath, S. L. Denham, L. Smith, H. Honing, A. Hazan, P. Holonowicz, H. Purwins
- 4:35pm **Modeling Natural Sounds with Modulation Cascade Processes, R. E. Turner, M. Sahani**
- 5:00pm *coffee break and poster viewing*
- 5:15pm **Dynamical Hierarchical Self-Organization of Harmonic, Motivic, and Pitch Categories,**
R. Marxer, P. Holonowicz, H. Purwins
- 5:25pm **What/When Causal Expectation Modelling in Monophonic Pitched and Percussive**
Audio, A. Hazan, P. Brossier, R. Marxer, H. Purwins
- 5:50pm **Modeling and Visualizing Tonality in North Indian Classical Music, P. Chordia, A. Rae**
- 6:00pm **Book-Adaptive and Book-Dependent Models to Accelerate Digitization of Early Music,**
L. Pugin, J. A. Burgoyne, D. Eck, I. Fujinaga
- 6:10pm *panel discussion*

Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives

Friday 7th - Saturday 8th December, 2007

Yael Niv, *Princeton University*

Matthew M. Botvinick, *Princeton University*

Andrew C. Barto, *University of Massachusetts, Amherst*

The aim of this workshop is to discuss current ideas from computer science, psychology and neuroscience regarding learning and control of hierarchically structured behavior. Psychological research has long emphasized that human behavior is hierarchically structured. Indeed, a hierarchical organization of human behavior that matches the hierarchical structure of real-world problems has been the focus of much empirical and theoretical research, and has played a pivotal role in research on organized, goal-directed behavior. Behavioral hierarchy has been of longstanding interest within neuroscience as well, where it has been considered to relate closely to prefrontal cortical function. The prefrontal cortex, which, with its high cognitive functions, remains the most poorly understood area of the brain, has been repeatedly implicated in supporting and executing hierarchical learning and control. In yet a third field, recent developments within machine learning have led to the emergence of 'hierarchical reinforcement learning'. This line of research has begun investigating in depth how optimal control can learn, and make use of, hierarchical structures, specifically, how hierarchies of skills (also termed options, macros or temporally abstract actions) could be learned and utilized optimally.

This workshop brings together front-line researchers from each of these fields, with the aim of gleaning new insights by integrating knowledge from these somewhat disparate areas of active research. The overarching goal is to facilitate sharing of ideas such as to potentially advance research in one field based on ideas and knowledge from other fields. The coming together of these three communities is especially exciting because, arguably, some of the most profound developments in psychology and neuroscience in the last two decades have stemmed from the use of normative ideas from reinforcement learning in thinking about and studying behavior and the brain. There is thus much promise in forging links between the long legacy of insight into human cognition and the more recent normative study of hierarchical control.

Hierarchical Organization of Behavior: Computational, Psychological and Neural Perspectives (Part 2)

Saturday December 8, 2007

Organizers: Yael Niv, Matthew M. Botvinick & Andrew C. Barto

Morning session: 9:00am–12:00pm

- 9:00am **Hierarchical organization of intelligence: Ethological and AI perspectives**, *J. Bryson*
- 9:25am **Structured event complexes in the human prefrontal cortex**, *J. Grafman & F. Kruger*
- 9:50am **Prefrontal topography of cognitive control according to levels of abstraction**, *K. Christoff*
- 10:15am **Panel Discussion**, *Bryson, Grafman, Kruger & Christoff*
- 10:35am *coffee break*
- 10:55am **Automatic induction of MAXQ hierarchies**,
N. Mehta, M. Wynkoop, S. Ray, P. Tadepalli & T. Diettrich
- 11:20am **Hierarchical lookahead agents: A preliminary report**, *B. Marthi, S. Russell & J. Wolfe*
- 11:45am **Panel Discussion**, *Dietterich & Wolfe*

Afternoon session: 3:00pm–6:00pm

- 3:00pm **The problem of decisiveness in adaptive behavior**, *R. Sutton*
- 3:25pm **Hierarchical apprenticeship learning with applications to quadruped locomotion**,
Z. Kolter, P. Abbeel & A. Ng
- 3:50pm **Reinforcement learning with multiple, qualitatively different state representations**,
H. van Seijen, B. Bakker & L. Kester
- 4:15pm **Panel Discussion**, *Sutton, Kolter & van Seijen*
- 4:35pm *coffee break*
- 4:55pm **Addressing the American problem by modeling cognitive development**, *Z. Stein*
- 5:20pm **Intrinsically motivated hierarchical reinforcement learning**, *A. Barto*
- 5:45pm **Panel Discussion**, *Stein & Barto*
- Summary & Wrap up*

NOTES

NOTES

NOTES

NOTES

NOTES