2013 WORKSHOP BOOK

NIPS 2013

NEURAL INFORMATION PROCESSING SYSTEMS WORKSHOP

TUTORIALS
December 5, 2013
Harrah's & Harveys
Lake Tahoe, Nevada

CONFERENCE SESSIONS
December 6 - 8, 2013
Harrah's & Harveys
Lake Tahoe, Nevada

WORKSHOPS
December 9 - 10, 2013
Harrah's & Harveys
Lake Tahoe, Nevada

Sponsored by the Neural Information Processing System Foundation, Inc

The technical program includes 6 invited talks and 360 accepted papers, selected from a total of 1420 submissions considered by the program committee. Because the conference stresses interdisciplinary interactions, there are no parallel sessions.

Papers presented at the conference will appear in "Advances in Neural Information Processing 26," edited by Léon Bottou, Chris Burges, Max Welling, Zoubin Ghahramani and Kilian Weinberger

Neural Information Processing Systems Foundation

TABLE OF CONTENTS

ORGANIZING COMMITTEE

General Chairs: **Léon Bottou** (Microsoft Research), **Chris J. C. Burges** (Microsoft Research)
Program Chairs: **Max Welling** (University of Amsterdam), **Zoubin Ghahramani** (Univ. of Cambridge)
Tutorials Chair: **Neil Lawrence** (University of Sheffield)
Workshop Chairs: **Rich Caruana** (Microsoft Research), **Gunnar Rätsch** (Memorial Sloan-Kettering Cancer Center)
Demonstration Chair: **Russ Salakhutdinov** (Univ. of Toronto)
Publications Chair and Electronic Proceedings Chair: **Kilian Weinberger** (Washington Univ)
Program Manager: **Hong Ge** (University of Cambridge)

PROGRAM COMMITTEE

Jacob Abernethy (University of Pennsylvania)
Ryan Adams (Harvard University)
Alekh Agarwal (Microsoft Research)
Cedric Archambeau (Xerox)
Francis Bach (INRIA-ENS)
Serge Belongie (UCSD)
Matthias Bethge (University of Tübingen)
Jeff Bilmes (University of Washington)
Karsten Borgwardt (MPI for Intelligent Systems)
Miguel Carreira-Perpiñan (UC Merced)
Gal Chechik (Bar Ilan University)
Corinna Cortes (Google)
Koby Crammer (Technion)
John Cunningham (Washington University in St Louis)
Ofer Dekel (Microsoft Research)
Emily Fox (University of Washington)
Mohammad Ghavamzadeh (INRIA)
Amir Globerson (Hebrew University of Jerusalem)
Dilan Görür (Yahoo Research)
Stefan Harmeling (MPI for Intelligent Systems)
Elad Hazan (Technion)
Tamir Hazan (TTI Chicago)
Matt Hoffman (Adobe)
Daniel Hsu (Microsoft Research)
Alex Ihler (UC Irvine)
Prateek Jain (Microsoft Research)
Dominik Janzing (MPI for Intelligent Systems)
Rong Jin (Michigan State University)
Andreas Krause (ETH Zürich)
Gert Lanckriet (UCSD)
Hugo Larochelle (Sherbrooke University)
Neil Lawrence (University of Sheffield)
Jure Leskovec (Stanford University)
Christina Leslie (Memorial Sloan-Kettering Cancer Center)
Fei Fei Li (Stanford University)
Han Liu (Princeton)
Mehryar Mohri (NYU and Google)
Claire Monteleoni (George Washington University)
Iain Murray (University of Edinburgh)
Guillaume Obozinski (Ecole des Ponts-ParisTech)
Jan Peters (TU Darmstadt)
Jonathan Pillow (University of Texas at Austin)
Massimiliano Pontil (UCL)
Gunnar Rätsch (Memorial Sloan-Kettering Cancer Center)
Maxim Raginsky (UIUC)
Deva Ramanan (UC Irvine)
Marc'Aurelio Ranzato (Google)
Lorenzo Rosasco (University of Genova)
Daniel Roy (University of Cambridge)
Cynthia Rudin (MIT)
Odelia Schwartz (Albert Einstein College of Medicine)
Aarti Singh (CMU)
Satinder Singh (University of Michigan)
Le Song (Georgia Tech)
David Sontag (NYU)
Suvrit Sra (MPI for Intelligent Systems)
Bharath Sriperumbudur (University of Cambridge)
Ambuj Tewari (University of Michigan)
Richard Turner (University of Cambridge)
Nuno Vasconcelos (UCSD)
S.V.N. Vishwanathan (Purdue University)
Frank Wood (University of Oxford)
Jennifer Wortman Vaughan (Microsoft Research)
Jieping Ye (Arizona State University)
Angela Yu (UCSD)
Jun Zhu (Tsinghua University)
Andrew Zisserman (University of Oxford)

CORE LOGISTICS TEAM

The organization and management of NIPS would not be possible without the help of many volunteers, students, researchers and administrators who donate their valuable time and energy to assist the conference in various ways. However, there is a core team at the Salk Institute whose tireless efforts make the conference run smoothly and efficiently every year. This year, NIPS would particularly like to acknowledge the exceptional work of:

Lee Campbell - IT Manager
Chris Hiestand - Webmaster
Ramona Marchand - Administrator
Mary Ellen Perry - Executive Director

2013 EXHIBITORS

Cambridge University Press
Robert Bosch LLC
CRC\Taylor & Francis Group
Now Publisher
The MIT Press
Springer
Pivotal

NIPS FOUNDATION OFFICERS & BOARD MEMBERS

President	Terrence Sejnowski, The Salk Institute
Treasurer	Marian Stewart Bartlett, University of California, San Diego
Secretary	Michael Mozer, University of Colorado, Boulder
Legal Advisor	Phil Sotel, Pasadena, CA
Executive	Peter Bartlett, Queensland University and University California, Berkeley Yoshua Bengio, University of Montreal John Lafferty, Carnegie Mellon University Fernando Pereira, Google Research John Shawe-Taylor, University College London Rich Zemel, University of Toronto Chris Williams, University of Edinburg
Emeritus Members	Gary Blasdel, Harvard Medical School T. L. Fine, Cornell University Eve Marder, Brandeis University
Advisory Board	Sue Becker, McMaster Univ, Ontario, Canada Jack Cowan, University of Chicago Thomas G. Dietterich, Oregon State Univ. Stephen Hanson, Rutgers University Michael I. Jordan, UC, Berkeley Michael Kearns, University of Pennsylvania Scott Kirkpatrick, Hebrew University, Daphne Koller, Stanford University Todd K. Leen, Oregon Graduate Institute Richard Lippmann, MIT Bartlett Mel, University of Southern California John Moody, ICSI, Berkeley and Portland John C. Platt, Microsoft Research Gerald Tesauro, IBM Watson Labs Sebastian Thrun, Stanford University Dave Touretzky, Carnegie Mellon University Lawrence Saul, UC San Diego Bernhard Schölkopf, Max Planck Institute Dale Schuurmans, Univ. of Alberta, Canada Sara A. Solla, Northwestern University Yair Weiss, Hebrew University of Jerusalem

MESSAGE FROM THE PRESIDENT

PHIL SOTEL

The NIPS community has lost a valued friend and a link to its founding fathers. Phil Sotel, who served as the pro bono general counsel of the NIPS Foundation since it was formed by Ed Posner at Caltech in 1992, died on October 6, 2013 after a brief and previously undiagnosed illness at the age of 77. He guided NIPS through its early years and had been a steady hand in advising the board as NIPS grew into a major international meeting. He attended annual NIPS Board meetings and advised the Foundation on major decisions. He put in place corporate procedures that have served the Foundation well. The road that we have traveled over the last 26 years would have had a lot more bumps without Phil's advice and the next 26 years will be more challenging without him. NIPS owes Ed Posner a debt of gratitude not only for founding NIPS, but also for convincing his friend Phil Sotel to oversee its growth and health.

Phil was involved with petroleum exploration in Indonesian and the Philippines and although he was based in Pasadena, he had a ranch in Colorado. He had close ties with Caltech and was on the wine committee at the Caltech Athenium. He was a member of the Hoover Institution at Stanford and the Pacific Council on International Policy, the West Coast affiliate of the Council on Foreign Relations. He had a broad perspective on the world, and many friends in commerce and academia, which helped NIPS grow into an organization that has had a significant impact on the development of the knowledge economy.

The first NIPS conference and workshop was held in Denver in 1987. It attracted a broad range of scientists and engineers with diverse backgrounds and interests eager to solve intractable problems using massively-parallel architectures and learning algorithms to deal with the high dimensionality of the parameters spaces and data sets. Brains were an inspiration and an existence proof that these problems could in fact be solved. NIPS served as an incubator for a field now widely known as machine learning, with deep roots in artificial intelligence, statistics, computational neuroscience, cognitive science, computer vision, control theory, speech recognition, neuromorphic engineering and many other disciplines that rely on the computational tools that have been developed by the NIPS community over the past 26 years. The annual NIPS conference and workshop has grown from 600 participants in 1987 to over 1,600 in 2012.

The Board of Trustees of the NIPS Foundation includes recent general chairs and also relies on an Advisory Board for continuity with past conferences. The Foundation oversees the annual meeting and handles its infrastructure, with the Board making decisions on where the meetings are held and who is asked to serve on organizing committees. The program committee chairs have the important task of choosing the best submissions and invited lecturers. Workshops are the crucibles for future advances that keep NIPS healthy. The NIPS Foundation thanks the members of all the organizing committees who have given their valuable time and service over the years.

Terry Sejnowski
La Jolla, CA, November 4, 2013

SPONSORS

NIPS gratefully acknowledges the generosity of those individuals and organizations who have provided financial support for the NIPS 2013 conference. The financial support enabled us to sponsor student travel and participation, the outstanding paper awards, the demonstration track and the opening buffet.

AFOSR continues to expand the horizon of scientific knowledge through its leadership and management of the Air Force's basic research program. As a vital component of the Air Force Research Laboratory (AFRL), AFOSR's mission is to support Air Force goals of control and maximum utilization of air, space, and cyberspace.

AFOSR accomplishes its mission by investing in basic research efforts for the Air Force in relevant scientific areas. Central to AFOSR's strategy is the transfer of the fruits of basic research to industry, the supplier of Air Force acquisitions; to the academic community which can lead the way to still more accomplishment; and to the other directorates of Air Force Research Laboratory (AFRL) that carry the responsibility for applied and development research leading to acquisition.

Microsoft Research is dedicated to pursuing innovation through basic and applied research in computer science and software engineering. Basic long-term research, unconstrained by the demands of product cycles, leads to new discoveries and lays the foundation for future technology breakthroughs that can define new paradigms, such as the current move toward cloud computing and software-plus-services. Applied research focuses on the near-term goal of improving products by transferring research findings and innovative technology to development teams. By balancing basic and applied research, and by maintaining an effective bridge between the two, Microsoft Research continually advances the state of the art in computer science and redefines the computing experience for millions of people worldwide. Microsoft Research has more than 1,100 scientists and engineers specializing in over 60 disciplines and includes some of the world's finest computer scientists, sociologists, psychologists, mathematicians, physicists, and engineers, working in our worldwide locations.

Amazon.com strives to be Earth's most customer-centric company where people can find and discover virtually anything they want to buy online. Amazon's evolution from Web site to e-commerce partner to development platform is driven by the spirit of innovation that is part of the company's DNA. The world's brightest technology minds come to Amazon.com to research and develop technology that improves the lives of shoppers, sellers and developers around the world. At Amazon, our Machine Learning team is comprised of technical leaders who develop planet-scale platforms for machine learning on the cloud, assist in the benchmarking and future development of existing machine learning applications across Amazon, and help develop novel and infinitely-scalable applications.

Google's mission is to organize the world's information and make it universally accessible and useful. Perhaps as remarkable as two Stanford research students having the ambition to found a company with such a lofty objective is the progress the company has made to that end. Ten years ago, Larry Page and Sergey Brin applied their research to an interesting problem and invented the world's most popular search engine. The same spirit holds true at Google today. The mission of research at Google is to deliver cutting-edge innovation that improves Google products and enriches the lives of all who use them. We publish innovation through industry standards, and our researchers are often helping to define not just today's products but also tomorrow's.

Helping over a billion people share and connect around the globe requires constant innovation. At Facebook, research permeates everything we do. Here, research is more than a lab—it's a way of doing things.

At Facebook, we believe that the most interesting academic problems are derived from real-world problems. Our researchers work on cutting edge research problems with a practical focus and push product boundaries every day. At the same time, they are publishing papers, giving talks, attending and hosting conferences and collaborating with the academic community. Our research teams are an integral part of the engineering organization and work with real user data to solve real-world problems that impact millions of people.

SKYTREE® THE MACHINE LEARNING COMPANY

Skytree®–The Machine Learning Company® is disrupting the Advanced Analytics market with a Machine Learning platform that gives organizations the power to discover deep analytic insights, predict future trends, make recommendations and reveal untapped markets and customers. Predictive Analytics is quickly becoming a must-have technology in the age of Big Data, and Skytree is at the forefront with enterprise-grade Machine Learning. Skytree's flagship product – Skytree Server – is the only general purpose scalable Machine Learning system on the market, built for the highest accuracy at unprecedented speed and scale.

United Technologies Research Center

United Technologies Research Center delivers the world's most advanced technologies, innovative thinking and disciplined research to the businesses of United Technologies -- industry leaders in aerospace propulsion, building infrastructure and services, heating and air conditioning, fire and security systems and power generation. Founded in 1929, UTRC is located in East Hartford, Connecticut (U.S.), with an office in Berkeley, California, and research and development centers in Shanghai, China, and Cork, Ireland. UTRC currently has several open roles for people with strong machine learning and distributed analytics skills to support service technologies across a wide array of applied industrial applications. If you're strong technically and enjoy working across a broad array of technical domains, UTRC may be the place for you.

IBM Research

IBM Research is a research and development organization consisting of twelve laboratories, worldwide. Major undertakings at IBM Research have included the invention of innovative materials and structures, high-performance microprocessors and computers, analytical methods and tools, algorithms, software architectures, methods for managing, searching and deriving meaning from data and in turning IBM's advanced services methodologies into reusable assets. IBM Research's numerous contributions to physical and computer sciences include the Scanning Tunneling Microscope and high temperature superconductivity, both of which were awarded the Nobel Prize. IBM Research was behind the inventions of the SABRE travel reservation system, the technology of laser eye surgery, magnetic storage, the relational database, UPC barcodes and Watson, the question-answering computing system that won a match against human champions on the Jeopardy! television quiz show. The Watson technology is now being commercialized as part of a project with healthcare company WellPoint. IBM Research is home to 5 Nobel Laureates, 9 US National Medals of Technology, 5 US National Medals of Science, 6 Turing Awards, and 13 Inductees in the National Inventors Hall of Fame.

TWO SIGMA

We are a technology company that applies a rigorous, scientific method-based approach to investment management. Since our founding in 2001, Two Sigma's vision has been to develop technological innovations that intelligently analyze the world's data to consistently deliver value for our clients. Our technology – inspired by a diverse set of fields including artificial intelligence and distributed computing – and our commitment to Research & Development aim to ensure that our methods are constantly improving and advancing.

D E Shaw & Co

Headquartered in New York City, the D. E. Shaw group is a global investment and technology development firm with offices in North America, Europe, and Asia. Since its organization in 1988 by a former Columbia University computer science professor, David E. Shaw, the firm has earned an international reputation for successful investing based on financial innovation, careful risk management, and the quality and depth of our staff. Our investment activities are based on both mathematical models and our staff's expertise, and our multi-disciplinary approach combines insights from quantitative fields, software development, sector expertise, and finance. We offer the benefits of being one of the world's largest, most established alternative investment managers, with a world-class technology infrastructure, deep research capabilities, and programs that facilitate the ongoing growth and internal mobility of staff. We have a long history of looking for candidates who aren't conventional "financial types," and our culture doesn't fit the typical corporate mold.

DRW TRADING GROUP

DRW Trading Group (DRW) is a principal trading organization. This means that all of our trading is for our own account and risk, and all of our methods, systems and applications are solely for our own use. Unlike hedge funds, brokerage firms and banks, DRW has no customers, clients, investors or third party funds. Our trading spans a wide range of asset classes, instruments, geographies and trading venues, with a focus on trading listed, centrally-cleared instruments.

Founded in 1992, our mission is to empower a team of exceptional individuals to identify and capture trading opportunities in the global markets by leveraging and integrating technology, risk management and quantitative research. With that spirit, DRW has embraced the integration of trading and technology by devoting extensive time, capital and resources to develop fast, precise and reliable infrastructure and applications. DRW has a flexible and entrepreneurial culture that cultivates creativity and practicality.

YAHOO! LABS

Yahoo Labs is the scientific engine powering one of the most trafficked Internet destinations worldwide. From idea to product innovation, Yahoo Labs is responsible for the algorithms behind the quality of the Web experience for hundreds of millions of users. We impact more than 800 million people in 60 countries who use Yahoo, and we do it from some of the most interesting, diverse, creative and inspiring locations on the planet. Our scientists collaborate with each other and with scientists outside Yahoo, pioneering innovations that improve the Yahoo experience in both evolutionary and revolutionary ways. Yahoo Labs scientists invent the technologies of the future, and then make them a reality today.

PDT PARTNERS

PDT Partners is a top quantitative hedge fund where world class researchers analyze rich data to develop and deploy model-driven algorithmic trading strategies. We offer a strong track record of hiring, challenging and retaining scientists interested in conducting research where the lab is the financial markets. Our researchers come from a variety of disciplines and backgrounds, having published in the top conferences and journals in machine learning, statistics, information theory, computational biology, pure and applied mathematics, theoretical and experimental physics, and operations research.

Composed of a tight-knit community of researchers, technologists, and business professionals, we strive to build one of the best quantitative trading firms in the world. Our investment success is driven by rigorous research, state-of-the-art technology, and keen focus on risk management and trade execution. We accomplish our goals by valuing depth and expertise, encouraging intellectual curiosity, and seeking constant innovation.

TOYOTA

Toyota Research Institute of North America (TRI-NA) was established in 2008 as a division of Toyota Technical Center (TTC) in Ann Arbor, MI. Toyota has been pursuing Sustainable Mobility, which addresses four key priorities: advanced technologies, urban environment, energy, and partnerships with government and academia.

Recently Toyota Motor Corporation (TMC) and its Lexus Division unveiled its advanced active safety research vehicle for the first time at the International CES to demonstrate ongoing efforts around autonomous vehicle safety technologies and explain Toyota's approach to reducing global traffic fatalities and injuries. The vehicle, based on a Lexus LS, advances the industry toward a new era of integrated safety management technologies (see 1).

The Lexus advanced active safety research vehicle is equipped with an array of sensors and automated control systems to observe, process and respond to the vehicle's surroundings. These include GPS, stereo cameras, radar and Light Detection and Ranging (LIDAR) laser tracking.

xerox®

Xerox Research Centre Europe research covers a broad spectrum of activities linked to information, data, documents and processes. The centre is internationally reputed for its expertise in computer vision, data analytics, natural language processing, machine learning, ethnography and process modelling.

The Machine Learning for Services group conducts fundamental and applied research in machine learning, computational statistics, and algorithmic mechanism design. Our research results are used in a wide range of applications, including relational learning, personalised content creation, large-scale recommender systems, and dynamic pricing.

The evidence-driven solutions we develop are part of Xerox services offerings. Xerox is the world leader in document management and business process outsourcing and research in Europe ensures that Xerox maintains that position.

Xerox Research Centre Europe is part of the global Xerox Innovation Group made up of 650 researchers and engineers in five world-class research centres. The Grenoble site is set in a park in the heart of the French Alps in a stunning location only a few kilometres from the city centre.

criteo

Criteo enables companies to engage and convert their customers online whether they are on a desktop, laptop, tablet or smartphone. Through its proprietary predictive algorithms, Criteo delivers performance- based online display advertising on real-time consumer data. Founded in 2005 in Paris, Criteo now employs more than 700 people across its 15 offices throughout the United States, Europe, Asia and Australia, serving more than 4,000 leading e-commerce companies across +35 countries globally. Our R&D team of 200+ engineers worldwide is building the next generation of digital advertising technologies that allow us to manage billions of ad impressions every month. For more information, please visit http://www.criteo.com

Springer
Machine Learning Journal

Springer Science+Business Media, LLC is an international e-publishing company for books and scholarly peer-reviewed journals in Science, Technology and Medicine (STM). E-content is distributed via numerous scientific databases, including SpringerLink, Springer Protocols, and SpringerImages. Book publications include major reference works, textbooks, monographs and various book series. Over 88,000 titles are available as e-books in 18 subject collections. The Machine Learning journal, published by Springer, is pleased to have supported the NIPS Best Student Awards.

PROGRAM HIGHLIGHTS

MONDAY, DECEMBER 9TH

Breakfast - 6:30 - 8:00 — See maps, Page 9

Registration — Harvey's 2nd Floor
7:00 - 11:00 AM & 3:30 - 6:30 PM

MONDAY WORKSHOPS

All workshops run from 7:30 - 10:30 AM &3:30 - 6:30 PM
Coffee breaks 9:00 - 9:30 AM and 5:00 - 5:30 PM

Advances in Machine Learning for Sensorimotor Control
Harvey's Emerald Bay 1

Big Learning : Advances in Algorithms and Data Management
Harvey's Emerald Bay B

Crowdsourcing: Theory, Algorithms and Applications
Harrah's Tahoe A+B

Deep Learning
Harrah's Sand Harbor II

Discrete Optimization in Machine Learning: Connecting Theory and Practice
Harrah's Sand Harbor I

Extreme Classification: Multi-Class & Multi-Label Learning with Millions of Categories
Harrah's Fallen+Marla

Frontiers of Network Analysis: Methods, Models, and Applications
Harvey's Emerald Bay 3

High-dimensional Statistical Inference in the Brain
Harvey's Emerald Bay 6

Large Scale Matrix Analysis and Inference
Harvey's Tallac

Machine Learning and Interpretation in Neuroimaging (Day 1)
Harvey's Sierra

Modern Nonparametric Methods in Machine Learning
Harvey's Zephyr

NIPS 2013 Workshop on Causality: Large-scale Experiment Design and Inference of Causal Mechanisms
Harrah's Glenbrook+Emerald

OPT2013: Optimization for Machine Learning
Harrah's Tahoe C

Output Representation Learning
Harrah's Sand Harbor III

Perturbations, Optimization, and Statistics
Harvey's Emerald Bay 2

Planning with Information Constraints for Control, Reinforcement Learning, Computational Neuroscience, Robotics and Games.
Harrah's Tahoe D

Probabilistic Models for Big Data
Harvey's Emerald Bay A

Randomized Methods for Machine Learning
Harvey's Emerald Bay 5

What Difference Does Personalization Make?
Harvey's Emerald Bay 4

TUESDAY, DECEMBER 10TH

Breakfast - 6:30 - 8:00 — See maps, Page 9

Registration — Harvey's 2nd Floor
7:00 - 11:00 AM

TUESDAY WOKSHOPS

All workshops run from 7:30 - 10:30 AM &3:30 - 6:30 PM
Coffee breaks 9:00 - 9:30 AM and 5:00 - 5:30 PM

Acquiring and Analyzing the Activity of Large Neural Ensembles
Harvey's Emerald Bay 4

Bayesian Optimization in Theory and Practice
Harvey's Emerald Bay A

Constructive Machine Learning
Harvey's Emerald Bay 1

Data Driven Education
Harrah's Tahoe D

Greedy Algorithms, Frank-Wolfe and Friends - A modern perspective
Harvey's Emerald Bay 6

Knowledge Extraction from Text (KET)
Harvey's Emerald Bay 2

Learning Faster From Easy Data
Harrah's Tahoe A

Machine Learning for Clinical Data Analysis and Healthcare
Harvey's Tallac

Machine Learning for Sustainability
Harrah's Glenbrook+Emerald

Machine Learning in Computational Biology
Harvey's Zephyr

Machine Learning Open Source Software: Towards Open Workflows
Harvey's Emerald Bay 5

Machine Learning and Interpretation in Neuroimaging (Day 2)
Harvey's Sierra

Neural Information Processing Scaled for Bioacoustics
Harrah's Tahoe C

New Directions in Transfer and Multi-Task: Learning Across Domains and Tasks
Harrah's Fallen+Marla

Resource-Efficient Machine Learning
Harvey's Emerald Bay 3

Topic Models: Computation, Application, and Evaluation
Harvey's Emerald Bay B

Workshop on Spectral Learning
Harrah's Tahoe B

Reception — Harrah's Sand Harbor
7:00 - 11:00

HARRAH'S & HARVEYS LOCATION MAPS

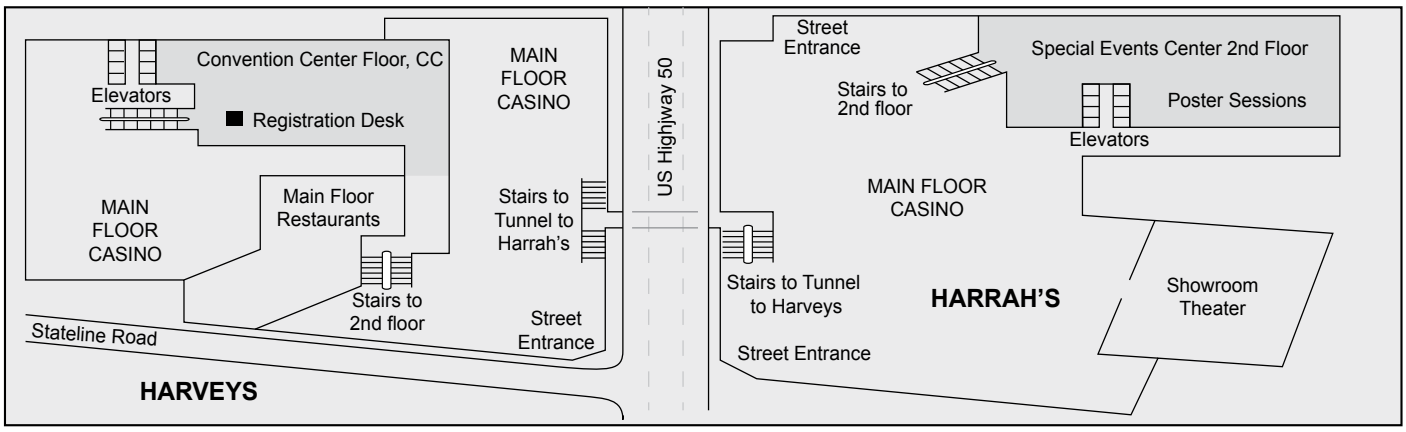

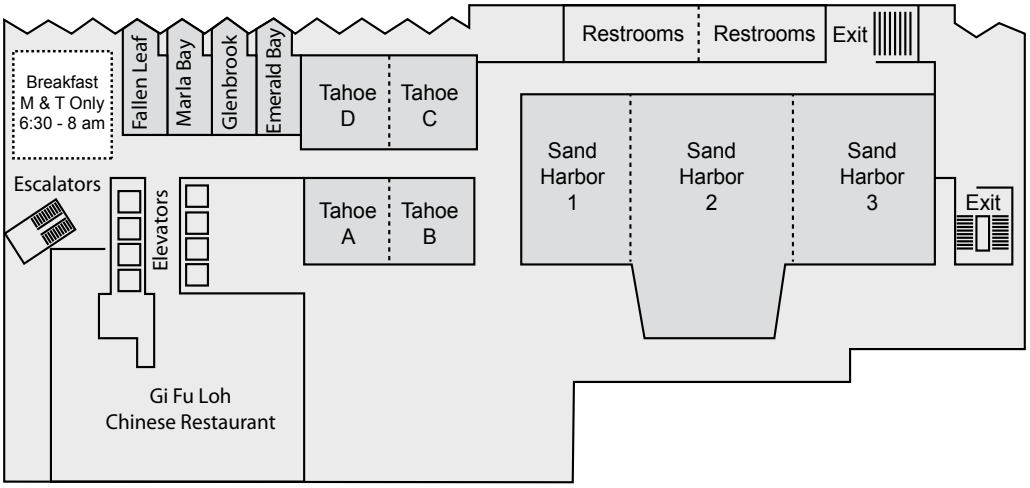

HARRAH'S SPECIAL EVENTS CENTER 2ND FLOOR

2 Breakfast Areas (dotted areas)

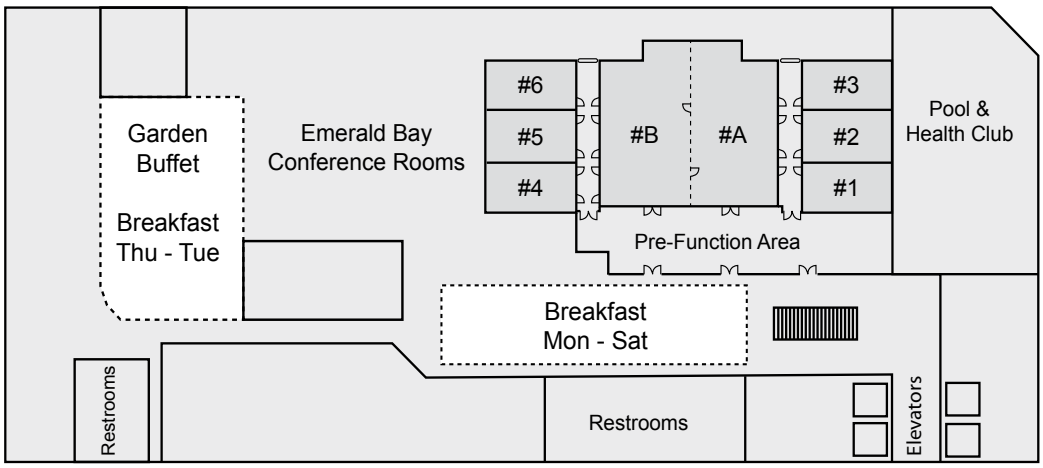

HARVEYS CONVENTION CENTER FLOOR, CC

HARVEYS

12TH FLOOR "TOP OF THE WHEEL"

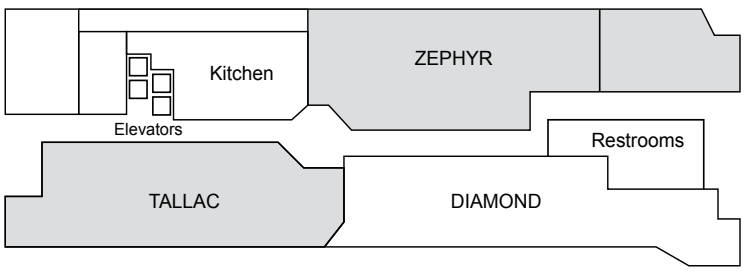

HARVEYS

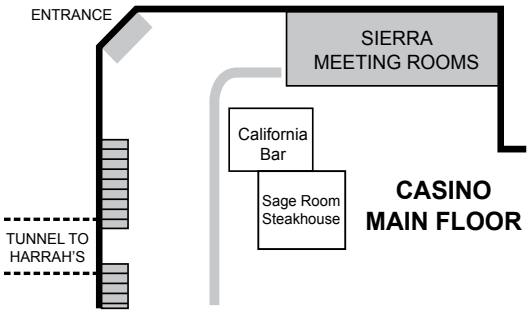

MONDAY WORKSHOPS

MONDAY WORKSHOP
7:30AM - 6:30PM

Advances in Machine Learning for Sensorimotor Control

ORGANIZERS:

Thomas Walsh	MIT
Alborz Geramifard	MIT
Jonathan How	MIT
Marc Deisenroth	Imperial College London
Jan Peters	TU Darmstadt

LOCATION:

Emerald Bay 1
Harveys Convention Center Floor (CC)

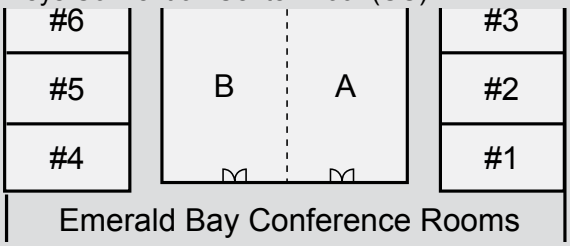

ABSTRACT:

Closed-loop control of systems based on sensor readings in uncertain domains is a hallmark of research in the Control, Artificial Intelligence, and Neuroscience communities. Various sensorimotor frameworks have been effective at controlling physical and biological systems, from flying airplanes to moving artificial limbs, but many techniques rely on accurate models or other concrete domain knowledge to derive useful policies. In systems where such specifications are not available, the task of generating usable models or even directly deriving controllers from data often falls in the purview of machine learning algorithms. Advances in machine learning, including non-parametric Bayesian modeling/ inference and reinforcement learning have increased the range, accuracy, and speed of deriving models and policies from data. However, incorporating modern machine learning techniques into real-world sensorimotor control systems can still be challenging due to the learner's underlying assumptions, the need to model uncertainty, and the scale of such problems. More specifically, many advanced machine learning algorithms rely either on strong distributional assumptions or random access to all possible data points, neither of which may be guaranteed when used with a specific control algorithm on a physical or biological system. In addition, planners need to consider, and learners need to indicate, uncertainty in the learned model/policy since some parameters may initially be uncertain but become known over time. Finally, most real-world sensorimotor control situations take place in continuous or high-dimensional environments and require real-time interaction, all of which are problematic for classical learning techniques. In order to overcome these difficulties, the modeling, learning, and planning components of a fully adaptive decision making system may need significant modifications. This workshop will bring together researchers from machine learning, control, and neuroscience that bridge this gap between effective planning and learning systems to produce better sensorimotor control. The workshop will be particularly concerned with the integration of machine learning and control components and the challenges of learning from limited data, modeling uncertainty, real-time execution, and the use of real-world data in complex sensorimotor environments. In addition to applications for mechanical systems, recent developments in biological motor control might be helpful to transfer to mechanical control systems and will also be a focus of the workshop. The workshop's domains of interest include a range of biological and physical systems with multiple sensors, including autonomous robots and vehicles, as well as complex real world systems, such as neural control, prosthetics, or healthcare where actions may take place over a longer time scale. High-level questions to be addressed (from a theoretical and practical perspective) include, but are not limited to: -How can we scale learning and planning techniques for the domain sizes encountered in real physical and biological systems? -How can online machine learning be used in high-frequency control of real-world systems? -How should planners use uncertainty measurements from approximate learned models for better exploration or to produce better plans in general? -How can successful supervised or unsupervised learning techniques be ported to sensorimotor control problems? - How can prior knowledge, including expert knowledge, user demonstrations, or distributional assumptions be incorporated into the learning/ planning framework? - How can safety and risk-sensitivity be incorporated into a planning/learning architecture? -How do biological systems deal with modeling, planning, and control under uncertainty? -How can we transfer biological insights to mechanical systems? -Do engineering insights have a biological explanation? - What lessons can be learned across disciplines between the control, neuroscience, and reinforcement learning communities, especially in their use of learning models? Website: http://acl.mit.edu/amlsc

WEBSITE:

http://acl.mit.edu/amlsc

MONDAY WORKSHOP
7:30AM - 6:30PM

Big Learning : Advances in Algorithms And Data Management

ORGANIZERS:

Xinghao Pan	UC Berkeley
Joseph Gonzalez	UC Berkeley
Haijie Gu	Carnegie Mellon University
Sameer Singh	University of Washington
Yucheng Low	GraphLab

LOCATION:

Emerald Bay B
Harveys Convention Center Floor (CC)

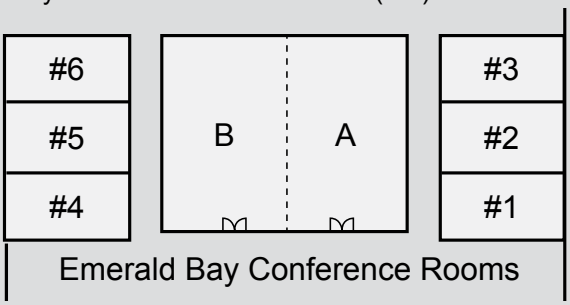

ABSTRACT:

Explosive growth in data and availability of cheap computing resources has sparked increasing interest in Big Learning within the Machine Learning community. Researchers are now taking on the challenge of parallelizing richly structured models with inherently serial dependencies and do not admit straightforward solutions. Database researchers, however, have a history of developing high performance systems that allow concurrent access while providing theoretical guarantees on correctness. In recent years, database systems have been developed specifically to tackle Big Learning tasks. This workshop aims to bring together the two communities and facilitate the cross-pollination of ideas. Rather than passively using DB systems, ML researchers can apply major DB concepts to their work; DB researchers stand to gain an understanding of the ML challenges and better guide the development of their Big Learning systems. The goals of the workshop are - Identify challenges faced by ML practitioners in Big Learning setting - Showcase recent and ongoing progress towards parallel ML algorithms - Highlight recent and significant DB research in addressing Big Learning problems - Introduce DB implementations of Big Learning systems, and the principle considerations and concepts underlying their designs Focal points for discussions and solicited submissions include but are not limited to: - Scalable data systems for Big Learning --- models and algorithms implemented, properties (availability, consistency, scalability, etc.), strengths and limitations - Distributed algorithms for online and batch learning - Parallel (multicore) algorithms for online and batch learning - Theoretical analysis of distributed and parallel learning algorithms - Implementation studies of large-scale distributed inference and learning algorithms --- challenges faced and lessons learnt Target audience includes industry and academic researchers from the various subfields relevant to large-scale machine learning, with a strong bias for either position talks that aim to induce discussion, or accessible overviews of the state-of-the-art.

WEBSITE:

http://biglearn.org/

Crowdsourcing: Theory, Algorithms And Applications

ORGANIZERS:

Qiang Liu	UC Irvine
Alex Ihler	UC Irvine
Dengyong Zhou	Microsoft Research
Jennifer Wortman Vaughan	Microsoft Research
Nikhil Devanur	Microsoft Research
Xi Chen	CMU

LOCATION:

Tahoe A+B
Harrah's Special Events Center 2nd Floor

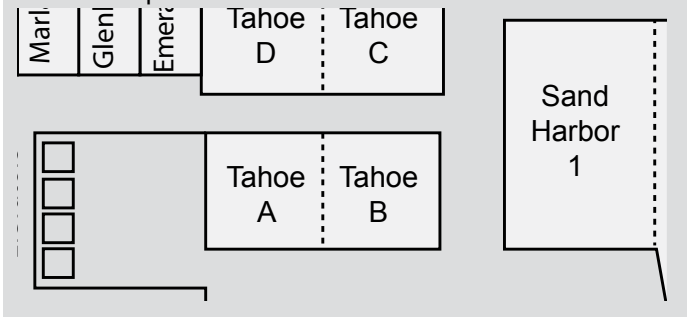

ABSTRACT:

All machine learning systems are an integration of data that store human or physical knowledge, and algorithms that discover knowledge patterns and make predictions to new instances. Even though most research attention has been focused on developing more efficient learning algorithms, it is the quality and amount of training data that predominately govern the performance of real-world systems. This is only amplified by the recent popularity of large scale and complicated learning systems such as deep networks, which require millions to billions of training data to perform well. Unfortunately, the traditional methods of collecting data from specialized workers are usually expensive and slow. In recent years, however, the situation has dramatically changed with the emergence of crowdsourcing, where huge amounts of labeled data are collected from large groups of (usually online) workers for low or no cost. Many machine learning tasks, such as computer vision and natural language processing are increasingly benefitting from data crowdsourced platforms such as Amazon Mechanical Turk and CrowdFlower. On the other hand, tools in machine learning, game theory and mechanism design can help to address many challenging problems in crowdsourcing systems, such as making them more reliable, efficient and less expensive. In this workshop, we call attention back to sources of data, discussing cheap and fast data collection methods based on crowdsourcing, and how it could impact subsequent machine learning stages. Furthermore, we will emphasize how the data sourcing paradigm interacts with the most recent emerging trends of machine learning in NIPS community. Examples of topics of potential interest in the workshop include (but are not limited to): Application of crowdsourcing to machine learning. Reliable crowdsourcing, e.g., label aggregation, quality control. Optimal budget allocation or active learning in crowdsourcing. Workflow design and answer aggregation for complex tasks (e.g., machine translation, proofreading). Pricing and incentives in crowdsourcing markets. Prediction markets / information markets and its connection to learning. Theoretical analysis for crowdsourcing algorithms, e.g., error rates and sample complexities for label aggregation and budget allocation algorithms.

WEBSITE:

http://www.ics.uci.edu/~qliu1/nips13_workshop/

MONDAY WORKSHOP
7:30AM - 6:30PM

Deep Learning

ORGANIZERS:

Yoshua Bengio	University of Montreal
Hugo Larochelle	Université de Sherbrooke
Russ Salakhutdinov	University of Toronto

LOCATION:

Sand Harbor 2
Harrah's Special Events Center 2nd Floor

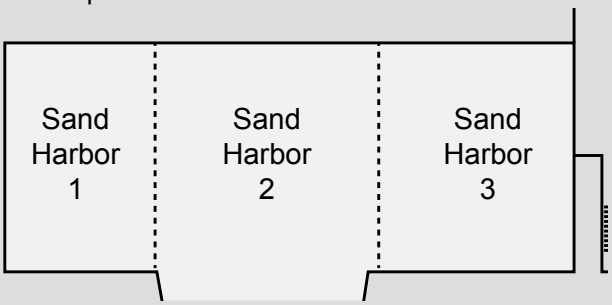

ABSTRACT:

Deep Learning algorithms attempt to discover good representations, at multiple levels of abstraction. There has been rapid progress in this area in recent years, both in terms of algorithms and in terms of applications, but many challenges remain. The workshop aims at bringing together researchers in that field and discussing these challenges, brainstorming about new solutions.

WEBSITE:

https://sites.google.com/site/deeplearningworkshopnips2013/

MONDAY WORKSHOP
7:30AM - 6:30PM

Discrete Optimization in Machine Learning: Connecting Theory and Practice

ORGANIZERS:

Stefanie Jegelka	UC Berkeley
Jeff Bilmes	University of Washington
Andreas Krause	ETH Zurich
Pradeep Ravikumar	UT Austin

LOCATION:

Sand Harbor 1
Harrah's Special Events Center 2nd Floor

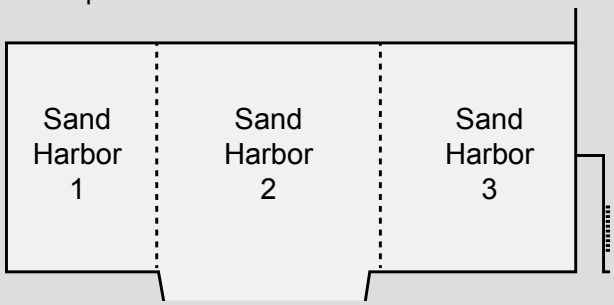

ABSTRACT:

We propose to hold a workshop that will bring together scientists interested in both discrete/combinatorial optimization and machine learning, and in doing so fostering interactions between researchers from different communities that typically do not have much chance to discuss ideas in person. There are many reasons why such a workshop will be beneficial to both communities. Solving optimization problems with ultimately discrete solutions is becoming increasingly important in machine learning. At the core of statistical machine learning is to make inferences from data, and when the variables underlying the data are discrete, both the tasks of inferring the model from data as well as performing predictions using the estimated model are inherently discrete optimization problems. Many of these optimization problems are notoriously hard. As a result, abundant and steadily increasing amounts of data -- despite being statistically beneficial -- quickly render standard off-the-shelf optimization procedures either impractical, intractable, or both. While many problems are hard in the worst case, the problems of practical interest are often much more well-behaved, or are well modeled by assuming properties that make them so. Indeed, many discrete problems in machine learning can possess beneficial structure; such structure has been an important ingredient in many successful (approximate) solution strategies. Examples include the marginal polytope, which is determined by the graph structure of the model, or sparsity that makes it possible to handle high dimensions. Symmetry and exchangeability are further exploitable characteristics. In addition, functional properties such as submodularity, a discrete analog of convexity, are proving to be useful to an increasing number of machine learning problems. One of the primary goals of this workshop is to provide a platform for exchange of ideas on how to discover, exploit, and deploy such structure. Machine learning, algorithms, discrete mathematics and combinatorics as well as applications in computer vision, speech, NLP, biology and network analysis are all active areas of research, each with an increasingly large body of foundational knowledge. This year, we want to ask questions that enable communication across such fields, and in particular between theory and practice, to wit: What are the most important structural properties arising in applications? What insights can we gain for these? How do theoretical analyses become important in practice? While discrete problems abound in harsh complexity results, many "heuristics" work very well in practice. What distinguishes these cases from worst cases? What theoretical insights or inspiration can be obtain from these empirical successes? How can we efficiently introduce novel realistic structural prior information? What are relevant techniques that the machine learning community should know about? How can we leverage the connection between machine learning and discrete optimization to gain a deeper understanding about all of the above questions?

WEBSITE:

http://discml.cc/

Extreme Classification: Multi-Class & Multi-Label Learning with Millions of Categories

ORGANIZERS:

Manik Varma — Microsoft Research India
John Langford — Microsoft Research NYC

LOCATION:

Fallen+Marla
Harrah's Special Events Center 2nd Floor

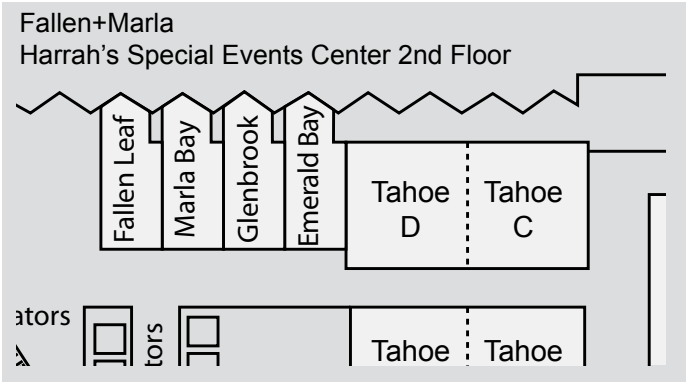

ABSTRACT:

Extreme classification, where one needs to deal with multi-class and multi-label problems involving a very large number of categories, has opened up a new research frontier in machine learning. Many challenging applications, such as photo and video annotation and web page categorization, can benefit from being formulated as supervised learning tasks with millions, or even billions, of categories. Extreme classification can also give a fresh perspective on core learning problems such as ranking and recommendation by reformulating them as multi-class/ label tasks where each item to be ranked or recommended is a separate category. Extreme classification raises a number of interesting research questions including those related to:

* Large scale learning and distributed and parallel training
* Efficient sub-linear prediction and prediction on a test-time budget
* Crowd sourcing and other efficient techniques for harvesting training data
* Dealing with training set biases and label noise
* Fine-grained classification
* Tackling label polysemy, synonymy and correlations
* Structured output prediction and multi-task learning
* Learning from highly imbalanced data
* Learning from very few data points per category
* Learning from missing and incorrect labels
* Feature extraction, feature sharing, lazy feature evaluation, etc.
* Performance evaluation
* Statistical analysis and generalization bounds

The workshop aims to bring together researchers interested in these areas to foster discussion and improve upon the state-of-the-art in extreme classification. Several leading researchers will present invited talks detailing the latest advances in the field. We also seek extended abstracts presenting work in progress which will be reviewed for acceptance as spotlight+poster or a talk. The workshop should be of interest to researchers in core supervised learning as well as application domains such as computer vision, computational advertising, information retrieval and natural language processing. We expect a healthy participation from both industry and academia.

WEBSITE:

http://research.microsoft.com/~manik/events/XC13/index.html

MONDAY WORKSHOP
7:30AM - 6:30PM

Frontiers of Network Analysis: Methods, Models, and Applications

ORGANIZERS:

Edo Airoldi	Harvard University
David Choi	Carnegie Mellon University
Aaron Clauset	University of Colorado
Khalid El-Arini	Facebook
Jure Leskovec	Stanford University

LOCATION:

Emerald Bay 3
Harveys Convention Center Floor (CC)

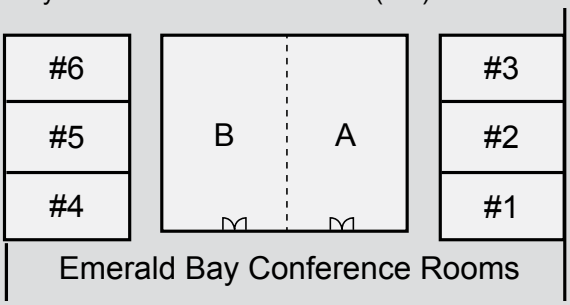

ABSTRACT:

Modern technology, including the World Wide Web, telecommunication devices and services, and large-scale data storage, has completely transformed the scale and concept of data in the sciences. Modern data sets are often enormous in size, detail, and heterogeneity, and are often best represented as highly annotated sequences of graphs. Although much progress has been made on developing rigorous tools for analyzing and modeling some types of large, complex, real-world networks, much work still remains and a principled, coherent framework remains elusive, in part because the analysis of networks is a growing and highly cross-disciplinary field. This workshop aims to bring together a diverse and cross-disciplinary set of researchers in order to both describe recent advances and to discuss future directions for developing new network methods in statistics and machine learning. By network methods, we broadly include those models and algorithms whose goal is to learn the patterns of interaction, flow of information, or propagation of effects in social, biological, and economic systems. We will also welcome empirical studies in applied domains such as the social sciences, biology, medicine, neuroscience, physics, finance, social media, and economics. While this research field is already broad and diverse, there are emerging signs of convergence, maturation, and increased methodological awareness. For example, in the study of information diffusion, social media and social network researchers are beginning to use rigorous tools to distinguish effects driven by social influence, homophily, or external processes -- subjects historically of intense interest amongst statisticians and social scientists. Similarly, there is a growing statistics literature developing learning approaches to study topics popularized earlier within the physics community, including clustering in graphs, network evolution, and random-graph models. Finally, learning methods are increasingly used in highly complex application domains, such as brain networks, and massive social networks like Facebook, and these applications are stimulating new scientific and practical questions that sometimes cut across disciplinary boundaries. Goals: The workshop's primary goal is to further facilitate the technical maturation of network analysis, promote greater technical sophistication and practical relevance, and identify future directions of research. To accomplish this, this workshop will bring together researchers from disciplines like computer science, statistics, physics, informatics, economics, sociology, with an emphasis on theoretical discussions of fundamental questions. The technical focus of the workshop is the statistical, methodological and computational issues that arise when modeling and analyzing large collections of heterogeneous and potentially dynamic network data. We seek to foster cross-disciplinary collaborations and intellectual exchange between the different communities and their respective ideas and tools. The communities identified above have long-standing interest in network modeling, and we aim to explore the similarities and differences both in methods and goals. The NIPS community serves as the perfect middle ground to enable effective communication of both applied and methodological concerns. We aim to once again bring together a diverse set of researchers to assess progress and stimulate further debate in an ongoing, open, cross-disciplinary dialogue. We believe this effort will ultimately result both in novel modeling approaches, and in the identification of new applications and open problems that may serve as guidance for future research directions. We welcome the following types of papers: 1. Research papers that introduce new models or apply established models to novel domains, 2. Research papers that explore theoretical and computational issues, or 3. Position papers that discuss shortcomings and desiderata of current approaches, or propose new directions for future research. We encourage authors to emphasize the role of learning and its relevance to the application domains at hand. In addition, we hope to identify current successes in the area, and will therefore consider papers that apply previously proposed models to novel domains and data sets.

WEBSITE:

http://snap.stanford.edu/networks2013/

MONDAY WORKSHOP
7:30AM - 6:30PM

High-dimensional Statistical Inference in the Brain

ORGANIZERS:

Alyson Fletcher	UC Santa Cruz
Dmitri Chklovskii	HHMI
Fritz Sommer	UC Berkeley
Ian Stevenson	University of Connecticut

LOCATION:

Emerald Bay 6
Harveys Convention Center Floor (CC)

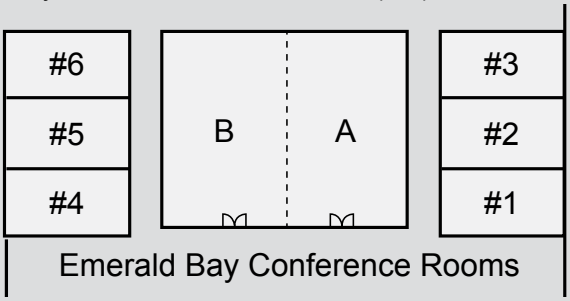

ABSTRACT:

Understanding high-dimensional phenomena is at the heart of many fundamental questions in neuroscience. How does the brain process sensory data? How can we model the encoding of the richness of the inputs, and how do these representations lead to perceptual capabilities and higher level cognitive function? Similarly, the brain itself is a vastly complex nonlinear, highly-interconnected network and neuroscience requires tractable, generalizable models for these inherently high-dimensional neural systems. Recent years have seen tremendous progress in high-dimensional statistics and methods for ``big data" that may shed light on these fundamental questions. This workshop seeks to leverage these advances and bring together researchers in mathematics, machine learning, computer science, statistics and neuroscience to explore the roles of dimensionality reduction and machine learning in neuroscience. Call for Papers We invite high quality submissions of extended abstracts on topics including, but not limited to, the following fundamental questions: -- How is high-dimensional sensory data encoded in neural systems? What insights can be gained from statistical methods in dimensionality reduction including sparse and overcomplete representations? How do we understand the apparent dimension expansion from thalamic to cortical representations from a machine learning and statistical perspective? -- What is the relation between perception and high-dimensional statistical inference? What are suitable statistical models for natural stimuli in vision and auditory systems? -- How does the brain learn such statistical models? What are the connections between unsupervised learning, latent variable methods, online learning and distributed algorithms? How do such statistical learning methods relate to and explain experience-driven plasticity and perceptual learning in neural systems? -- How can we best build meaningful, generalizable models of the brain with predictive value? How can machine learning be leveraged toward better design of functional brain models when data is limited or missing? What role can graphical models coupled with newer techniques for structured sparsity play in this dimensionality reduction? -- What are the roles of statistical inference in the formation and retrieval of memories in the brain? We wish to invite discussion on the very open questions of multi-disciplinary interest: for memory storage, how does the brain decode the strength and pattern of synaptic connections? Is it reasonable to conjecture the use of message passing algorithms as a model? -- Which estimation algorithms can be used for inferring nonlinear and inter-connected structure of these systems? Can new compressed sensing techniques be exploited? How can we model and identify dynamical aspects and temporal responses? We have invited researchers from a wide range of disciplines in electrical engineering, psychology, statistics, applied physics, machine learning and neuroscience with the goals of fostering interdisciplinary insights. We hope that active discussions between these groups can set in motion new collaborations and facilitate future breakthroughs on fundamental research problems.
Email: hdnips2013@rctn.org

WEBSITE:

http://users.soe.ucsc.edu/~afletcher/hdnips2013.html

Large Scale Matrix Analysis and Inference

ORGANIZERS:

Reza Bosagh Zadeh	Stanford University
Gunnar Carlsson	Stanford University
Michael Mahoney	Stanford University
Manfred Warmuth	Univ. of Calif. at Santa Cruz
Wouter Koolen	Queensland Univ. of Technology

LOCATION:

Tallac
Harveys 12th Floor “Top Of The Wheel”

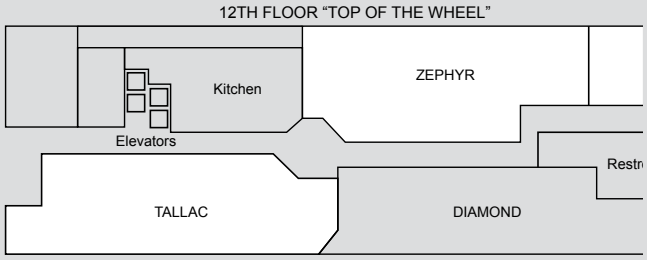

ABSTRACT:

RMuch of Machine Learning is based on Linear Algebra. Often, the prediction is a function of a dot product between the parameter vector and the feature vector. This essentially assumes some kind of independence between the features. In contrast matrix parameters can be used to learn interrelations between features: The (i,j)th entry of the parameter matrix represents how feature i is related to feature j. This richer modeling has become very popular. In some applications, like PCA and collaborative filtering, the explicit goal is inference of a matrix parameter. Yet in others, like direction learning and topic modeling, the matrix parameter instead pops up in the algorithms as the natural tool to represent uncertainty. The emergence of large matrices in many applications has brought with it a slew of new algorithms and tools. Over the past few years, matrix analysis and numerical linear algebra on large matrices has become a thriving field. Also manipulating such large matrices makes it necessary to to think about computer systems issues. This workshop aims to bring closer researchers in large scale machine learning and large scale numerical linear algebra to foster cross-talk between the two fields. The goal is to encourage machine learning researchers to work on numerical linear algebra problems, to inform machine learning researchers about new developments on large scale matrix analysis, and to identify unique challenges and opportunities. The workshop will conclude with a session of contributed posters.

WEBSITE:

http://largematrix.org

MLINI-13: Machine Learning And Interpretation In Neuroimaging

ORGANIZERS:

Irina Rish	IBM T.J. Watson Research Ctr
Georg Langs	Medical University of Vienna
Brian Murphy	Carnegie Mellon University
Kai-min Chang	Carnegie Mellon University
Guillermo Cecchi	IBM Research

LOCATION:

Sierra
Harveys Casino Main Floor

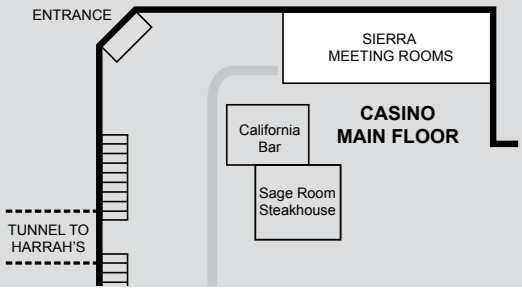

ABSTRACT:

This is a two-day workshop on the topic of machine learning approaches in neuroscience, neuroimaging, with a specific extension to behavioral experiments and psychology. We believe that machine learning has a prominent role in shaping how questions in neuroscience are framed, and that the machine-learning mind set is now entering modern psychology and behavioral studies. The goal of the workshop is to pinpoint the most pressing issues and common challenges across the neuroscience and machine learning, and to sketch future directions and open questions in the light of novel methodology. The workshop is aimed at offering a forum that joins machine learning, neuroscience, and psychology community, and should facilitate formulating and discussing the issues at their interface.

Motivated by two previous workshops, MLINI '11 and MLINI'12, we will center this workshop around invited talks, and two panel discussions. Triggered by these discussions, this year we plan to adapt the workshop topics to a less traditional scope that investigates the role of machine learning in neuroimaging of both animals and humans, as well as in behavioral models and psychology.

We intend to focus on the implications that follow from adopting multivariate machine-learning methods for studying brain function. In particular, we would like to investigate how these methods may be used to represent cognitive states, and what ramifications this has for consequent theories of cognition. Besides providing a rationale for the use of machine-learning methods in studying brain function, a further goal of this workshop is to identify shortcomings of state-of-the-art approaches and initiate research efforts that increase the impact of machine learning on cognitive neuroscience. Possible topics for discussion include the following main themes: I) machine learning and pattern recognition methodology in brain research, II) functional connectivity and dynamical models of brain activity, III) multi-modal analysis including mental state inference from behavioral data, and IV) linking machine learning, neuroimaging and neuroscience

WEBSITE:

https://sites.google.com/site/mlininips2013/

Modern Nonparametric Methods in Machine Learning

ORGANIZERS:

Arthur Gretton	UCL
Mladen Kolar	University of Chicago
John Lafferty	University of Chicago
Samory Kpotufe	TTI Chicago
Han Liu	Princeton University

LOCATION:

Zephyr
Harveys 12th Floor “Top Of The Wheel”

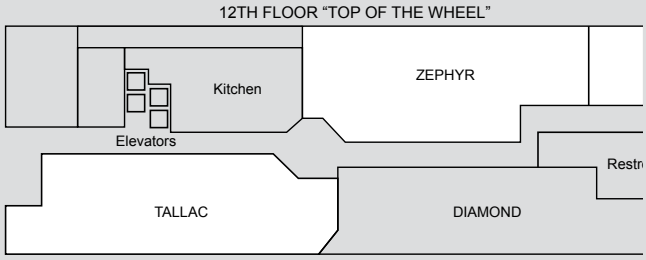

ABSTRACT:

Modern data acquisition routinely produces massive and complex datasets. Examples are data from high throughput genomic experiments, climate data from worldwide data centers, robotic control data collected overtime in adversarial settings, user-behavior data from social networks, user preferences on online markets, and so forth. Modern pattern recognition problems arising in such disciplines are characterized by large data sizes, large number of observed variables, and increased pattern complexity. Therefore, nonparametric methods which can handle generally complex patterns are ever more relevant for modern data analysis. However, the larger data sizes and number of variables constitute new challenges for nonparametric methods in general. The aim of this workshop is to bring together both theoretical and applied researchers to discuss these modern challenges in detail, share insight on existing solutions, and lay out some of the important future directions. Through a number of invited and contributed talks and a focused panel discussion, we plan to emphasize the importance of nonparametric methods and present challenges for modern nonparametric methods. In particular, we focus on the following aspect of nonparametric methods: A. General motivations for nonparametric methods: * the abundance of modern applications where little is known about data generating mechanisms (e.g., robotics, biology, social networks, recommendation systems) * the ability of nonparametric analysis to capture general aspects of learning such as bias-variance tradeoffs, and thus yielding general insight on the inherent complexity of various learning tasks. B. Modern challenges for nonparametric methods: * handling big data: while large data sizes are a blessing w.r.t. generalization performance, they also present a modern challenge for nonparametric learning w.r.t. time-efficiency. In this context, we need to characterize trade-off between time and accuracy, create online or stream-based solutions, and develop approximation methods. * larger problem complexity: large data is often paired with (1) large data dimension (number of observed variables), and (2) more complex target model spaces (e.g. less smooth regression function). To handle large data dimensions, likely solutions are methods that perform nonlinear dimension reduction, nonparametric variable selection, or adapt to the intrinsic dimension of the data. To handle the increased complexity of target model spaces, we require modern model selection procedures that can efficiently scale to modern data sizes while adapting to the complexity of the problem at hand.

WEBSITE:

http://sites.google.com/site/nips2012modernnonparametric

MONDAY WORKSHOP
7:30AM - 6:30PM

NIPS 2013 Workshop on Causality: Large-scale Experiment Design & Inference of Causal Mechanisms

ORGANIZERS:

Isabelle Guyon	Clopinet
Leon Bottou	Microsoft Research
Evelyne Viegas	Microsoft Research
Bernhard Schölkopf	MPI Tübingen
Alexander Statnikov	New York University

LOCATION:

Glenbrook+Emerald
Harrah's Special Events Center 2nd Floor

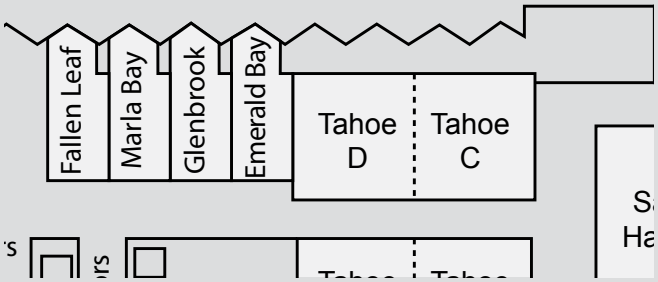

ABSTRACT:

The problem of attributing causes to effects is pervasive in science, medicine, economy and almost every aspects of our everyday life involving human reasoning and decision making. What affects your health? the economy? climate changes? The gold standard to establish causal relationships is to perform randomized controlled experiments. However, experiments are costly while non-experimental "observational" data collected routinely around the world are readily available. Unraveling potential cause-effect relationships from such observational data could save a lot of time and effort by allowing us to prioritize confirmatory experiments. This could be complemented by new strategies of incremental experimental design combining observational and experimental data.

The goal of this workshop is to discuss new methods of large scale experiment design and their application to the inference of causal mechanisms and promote their evaluation via a series of challenges. Emphasis will be put on capitalizing on massive amounts of available observational data to cut down the number of experiments needed, pseudo- or quasi-experiments, iterative designs, and the on-line acquisition of data with minimal perturbation of the system under study. The results of the cause-effect pairs challenge http://www.causality.inf.ethz.ch/cause-effect.php will be discussed.

WEBSITE:

http://clopinet.com/isabelle/Projects/NIPS2013/

OPT2013: Optimization for Machine Learning

ORGANIZERS:

Suvrit Sra — Max Planck ISI & CMU
Alekh Agarwal — Microsoft Research

LOCATION:

Tahoe C
Harrah's Special Events Center 2nd Floor

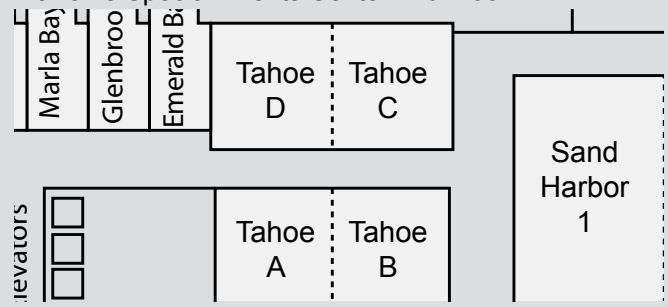

ABSTRACT:

OPT 2013 stands on significant precedent established by OPT 2008--OPT 2012 which were all very well-received NIPS workshops. The previous OPT workshops enjoyed packed (to overpacked) attendance, and this enthusiastic reception underscores the strong interest, relevance, and importance enjoyed by optimization in the ML community. This interest has grown remarkably strongly every year, no wonder, since optimization lies at the heart of most ML algorithms. Although classical textbook algorithms might sometimes suffice, the majority of ML problems require tailored methods based on a deeper understanding of learning task. Indeed, ML applications and researchers are driving some of the most cutting-edge developments in optimization today. This intimate relation of optimization with ML is the key motivation for our workshop, which aims to foster discussion, discovery, and dissemination of the state-of-the-art in optimization as relevant to machine learning.

I. INTRODUCTION
OPT workshops have previously covered a variety of topics, such as frameworks for convex programs (D. Bertsekas), the intersection of ML and optimization, especially SVM training (S. Wright), large-scale learning via stochastic gradient methods and its tradeoffs (L. Bottou, N. Srebro), exploitation of structured sparsity (Vandenberghe), randomized methods for extremely large-scale convex optimization (A. Nemirovski), complexity theoretic foundations of convex optimization (Y. Nesterov), distributed large-scale optimization (S. Boyd), asynchronous and sparsity based stochastic gradient (B. Recht), and algebraic techniques in machine learning (P. Parillo) among others. Several ideas propounded in these talks have become important research topics in ML and optimization --- especially in the field of randomized algorithms and stochastic gradient methods. An edited book "Optimization for Machine Learning" (S. Sra, S. Nowozin, and S. Wright; MIT Press, 2011) grew out of the first three OPT workshops, and contains high-quality contributions from many of the speakers and attendees. Much of the recent focus has been on stochastic methods, which are particularly suited to large and high-dimensional data sets. Theoretical and practical advances in these methods remain a topic of core interest to the workshop. The need for an up-to-date analysis of algorithms for nonconvex problems remains an important practical issue, whose importance becomes even more pronounced as ML tackles more and more complex mathematical models. This has particularly come to the fore given the recent successes of deep learning and dictionary learning methods. We also do not wish to ignore the not particularly large scale setting, where one does have time to wield substantial computational resources. In this setting, high-accuracy solutions and deep understanding of the lessons contained in the data are needed. Examples valuable to MLers may be exploration of genetic and environmental data to identify risk factors for disease; or problems dealing with setups where the amount of observed data is not huge, but the mathematical model is complex. Consequently, we encourage optimization methods on manifolds, ML problems with differential geometric antecedents, those using advanced algebraic techniques, and computational topology, for instance.

WEBSITE:

MONDAY WORKSHOP
7:30AM - 6:30PM

Output Representation Learning

ORGANIZERS:

Yuhong Guo	Temple University
Dale Schuurmans	University of Alberta
Richard Zemel	University of Toronto

LOCATION:

Sand Harbor III
Harrah's Special Events Center 2nd Floor

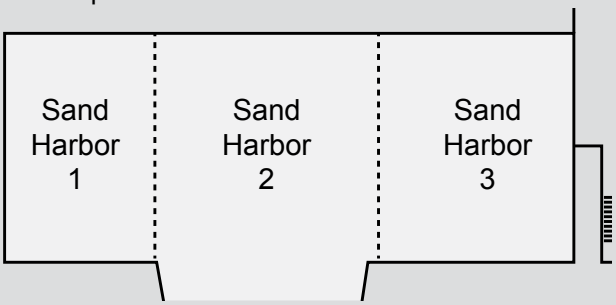

ABSTRACT:

Modern data analysis is increasingly facing prediction problems that have complex and high dimensional output spaces. For example, document tagging problems regularly consider large (and sometimes hierarchical) sets of output tags; image tagging problems regularly consider tens of thousands of possible output labels; natural language processing tasks have always considered complex output spaces. In such complex and high dimensional output spaces the candidate labels are often too specialized---leading to sparse data for individual labels---or too generalized---leading to complex prediction maps being required. In such cases, it is essential to identify an alternative output representation that can provide latent output categories that abstract overly specialized labels, specialize overly abstract labels, or reveal the latent dependence between labels. There is a growing body of work on learning output representations, distinct from current work on learning input representations. For example, in machine learning, work on multi-label learning, and particularly output dimensionality reduction in high dimensional label spaces, has begun to address the specialized label problem, while work on output kernel learning has begun to address the abstracted label problem. In computer vision, work on image categorization and tagging has begun to investigate simple forms of latent output representation learning to cope with abstract semantic labels and large label sets. In speech recognition, dimensionality reduction has been used to identify abstracted outputs, while hidden CRFs have been used to identify specialized latent outputs. In information retrieval and natural language processing, discovering latent output specializations in complex domains has been an ongoing research topic for the past half decade. The aim of this workshop is to bring these relevant research communities together to identify fundamental strategies, highlight differences, and identify the prospects for developing a set of systematic theory and methods for output representation learning. The target communities include researchers working on image tagging, document categorization, natural language processing, large vocabulary speech recognition, deep learning, latent variable modeling, and large scale multi-label learning.

WEBSITE:

https://sites.google.com/site/outputrepresentlearn2013/

MONDAY WORKSHOP
7:30AM - 6:30PM

Perturbations, Optimization, and Statistics

ORGANIZERS:

Tamir Hazan	University of Haifa
George Papandreou	Toyota Technological Institute Chicago
Sasha Rakhlin	University of Pennsylvania
Daniel Tarlow	Microsoft Research Cambridge

LOCATION:

Emerald Bay 2
Harveys Convention Center Floor (CC)

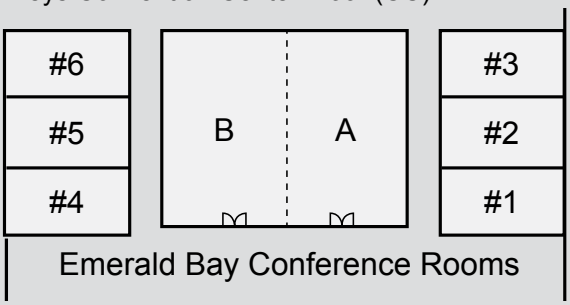

ABSTRACT:

In nearly all machine learning tasks, decisions must be made given current knowledge (e.g., choose which label to predict). Perhaps surprisingly, always making the best decision is not always the best strategy, particularly while learning. Recently, there is an emerging body of work on learning under different rules that apply perturbations to the decision procedure. These works provide simple and efficient learning rules with improved theoretical guarantees. This workshop will bring together the growing community of researchers interested in different aspects of this area, and it will broaden our understanding of why and how perturbation methods can be useful. Last year, at the highly successful NIPS workshop on Perturbations, Optimization, and Statistics, we looked at how injecting perturbations (whether it be random or adversarial "noise") into learning and inference procedures can be beneficial. The focus was on two angles: first, on how stochastic perturbations can be used to construct new types of probability models for structured data; and second, how deterministic perturbations affect the regularization and the generalization properties of learning algorithms. The goal of this workshop is to expand the scope of last year and also explore different ways to apply perturbations within optimization and statistics to enhance and improve machine learning approaches. This year, we would like to: (a) Look at exciting new developments related to the above core themes. (b) Emphasize their implications on topics that received less coverage last year, specifically highlighting connections to decision theory, risk analysis, game theory, and economics. More generally, we shall specifically be interested in understanding the following issues: * Repeated games and online learning: How to use random perturbations to explore unseens options in repeated games? How to exploit connections to Bayesian risk? * Adversarial Uncertainty: How to play complex games with adversarial uncertainty? What are the computational qualities of such solutions, and do Nash-equilibria exists in these cases? * Stochastic risk: How to average predictions with random perturbations to get improved generalization guarantees? How stochastic perturbations imply approximated Bayesian risk and regularization? * Dropout: How stochastic dropout regularizes learning of complex models and what is its generalization power? What are the relationships between stochastic and adversarial dropouts? * Robust optimization: In what ways can learning be improved by perturbing the input measurements? * Choice theory: What is the best way to use perturbations to compensate lack of knowledge? What lessons in modeling can machine learning take from random utility theory? * Theory: How does the maximum of a random process relate to its complexity? How can the maximum of random perturbations be used to measure the uncertainty of a system?

WEBSITE:

http://www.cs.toronto.edu/~dtarlow/pos13/

MONDAY WORKSHOP
7:30AM - 6:30PM

Planning with Information Constraints for Control, Reinforcement Learning, Computational Neuroscience, Robotics and Games

ORGANIZERS:

Hilbert Kappen	Radboud University
Naftali Tishby	Hebrew University Jerusalem
Jan Peters	TU Darmstadt
Evangelos Theodorou	Georgia Institute of Technology
David Wolpert	Santa Fe Institute
Pedro Ortega	University of Pennsylvania

LOCATION:

Tahoe D
Harrah's Special Events Center 2nd Floor

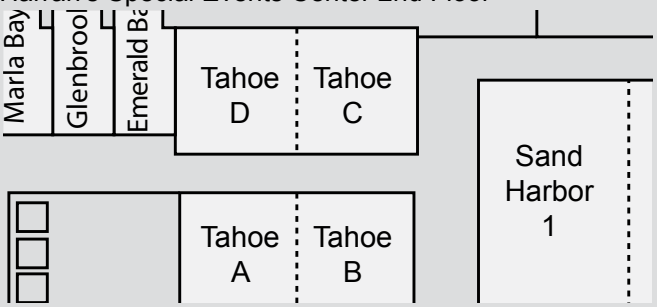

ABSTRACT:

How do you make decisions when there are way more possibilities than you can analyze? How do you decide under such information constraints? Planning and decision-making with information constraints is at the heart of adaptive control, reinforcement learning, robotic path planning, experimental design, active learning, computational neuroscience and games. In most real-world problems, perfect planning is either impossible (computational intractability, lack of information, diminished control) or sometimes even undesirable (distrust, risk sensitivity, level of cooperation of the others). Recent developments have shown that a single method, based on the free energy functional borrowed from thermodynamics, provides a principled way of designing systems with information constraints that parallels Bayesian inference. This single method -known in the literature under various labels such as KL-control, path integral control, linearly-solvable stochastic control, information-theoretic bounded rationality- is proving itself very general and powerful as a foundation for a novel class of probabilistic planning problems. The goal of this workshop is twofold: 1) Give a comprehensive introduction to planning with information constraints targeted to a wide audience with machine learning background. Invited speakers will give an overview of the theoretical results and talk about their experience in applications to control, reinforcement learning, computational neuroscience and robotics. 2) Bring together the leading researchers in the field to discuss, compare and unify their approaches, while interacting with the audience. Recent advances will be presented in a poster session based on contributed material. Furthermore, ample space will be given to state open questions and to sketch future directions.

WEBSITE:

http://www.seas.upenn.edu/~ope/workshop/

Probabilistic Models for Big Data

ORGANIZERS:

Neil Lawrence	University of Sheffield
Joaquin Quiñonero Candela	Facebook
Tianshi Gao	Facebook
James Hensman	University of Sheffield

LOCATION:

Emerald Bay A
Harveys Convention Center Floor (CC)

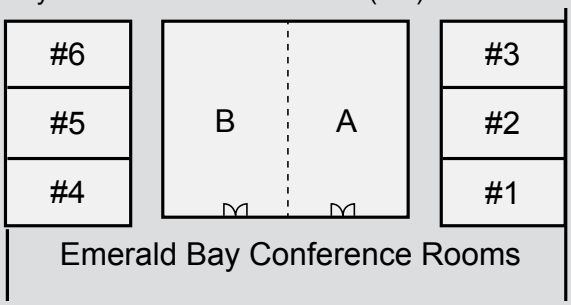

ABSTRACT:

Processing of web scale data sets has proven its worth in a range of applications, from ad-click prediction to large recommender systems. In most cases, learning needs to happen real-time, and the latency allowance for predictions is restrictive. Probabilistic predictions are critical in practice on web applications because optimizing the user experience requires being able to compute the expected utilities of mutually exclusive pieces of content. The quality of the knowledge extracted from the information available is restricted by complexity of the model. One framework that enables complex modelling of data is probabilistic modelling. However, its applicability to big data is restricted by the difficulties of inference in complex probabilistic models, and by computational constraints. This workshop will focus on applying probabilistic models to big data. Of interest will be algorithms that allow for inference in probabilistic models for big data such as stochastic variational inference and stochastic Monte Carlo. A particular focus will be on existing applications in big data and future applications that would benefit from such approaches. This workshop brings together leading academic and industrial researchers in probabilistic modelling and large scale data sets.

WEBSITE:

https://sites.google.com/site/probabilisticmodelsforbigdata/

Randomized Methods for Machine Learning

ORGANIZERS:

David Lopez-Paz — MPI, University of Cambridge
Quoc V. Le — CMU and Google
Alex Smola — CMU and Google

LOCATION:

Emerald Bay 5
Harveys Convention Center Floor (CC)

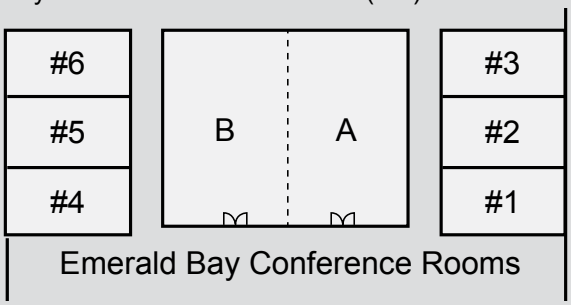

ABSTRACT:

As we enter the era of “big-data”, Machine Learning algorithms that resort in heavy optimization routines rapidly become prohibitive. Perhaps surprisingly, randomization (Raghavan and Motwani, 1995) arises as a computationally cheaper, simpler alternative to optimization that in many cases leads to smaller and faster models with little or no loss in performance. Although randomized algorithms date back to the probabilistic method (Erdős, 1947, Alon & Spencer, 2000), these techniques only recently started finding their way into Machine Learning. The most notable exceptions being stochastic methods for optimization, and Markov Chain Monte Carlo methods, both of which have become well-established in the past two decades. This workshop aims to accelerate this process by bringing together researchers in this area and exposing them to recent developments. The targeted audience are researchers and practitioners looking for scalable, compact and fast solutions to learn in the large-scale setting. Specific questions of interest include, but are not limited to: - Randomized projections: locality sensitive hashing, hash kernels, counter braids, count sketches, optimization. - Randomized function classes: Fourier features, Random Kitchen Sinks, Nystrom methods, Fastfood, Random Basis Neural networks. - Sparse reconstructions: compressed sensing, error correcting output codes, reductions of inference problems to binary. - Compressive approximations: min-hash, shingles, Bloom filters, coresets, random subsampling from streams. - Randomized dependence measures, component analysis, dimensionality reduction. - Extensions to less exploited tasks: density estimation, multitask and semi-supervised learning, deep and hierarchical models, feature learning, control, causality. - Hybrid strategies that combine optimization and randomization. - Sampling algorithms for Bayesian inference. - Random matrices and graphs. This one day workshop will feature invited tutorials and contributed short talks. Poster sessions, coffee breaks and a closing panel will encourage discussion between the attendants. We plan to collect a tightly edited collection of papers from the workshop in the form of a special issue or a book. This will allow faster dissemination of randomized methods in machine learning.

More information will be available at the official website www.randomizedmethods.org.

WEBSITE:

http://www.randomizedmethods.org

What Difference Does Personalization Make?

ORGANIZERS:

Dilan Gorur	Yahoo! Labs
Romer Rosales	LinkedIn
Olivier Chapelle	Criteo
Dorota Glowacka	University of Helsinki

LOCATION:

Emerald Bay 4
Harveys Convention Center Floor (CC)

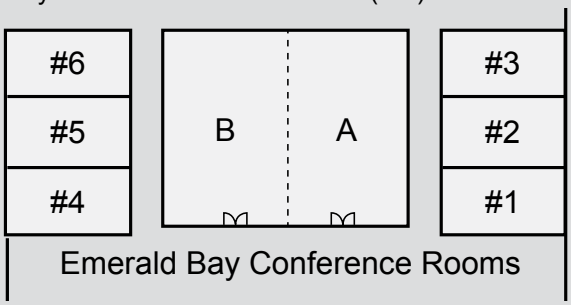

ABSTRACT:

Personalization has become an important research topic in machine learning fueled in part by its major significance in e-comerce and other businesses that try to tailor to user-specific preferences. Online products, news, search, media, and advertisement are some of the areas that have depended on some form of personalization to improve user satisfaction or business goals in general. In order to address personalization problems machine learning has long relied on tools such as collaborative filtering (matrix factorization) and models originally developed not necessarily for personalization. However, even though the data available for personalization has grown in richness and size, and the available processing power has also increased, the basic tenet for the methods used has not changed in a major way. It is possible that personalization requires a change in perspective, to learning the finer, user specific details in the data. It may be necessary to develop modeling and evaluation approaches different than those developed for more general purposes. We aim to motivate these and new discussions to foster innovation in the area of machine learning for personalization. Research efforts on this topic outside of the NIPS community could provide useful insights into developing new methods and points of view. This workshop will bring together experts in various fields including machine learning, data mining, information retrieval and social sciences, with the goal of understanding the current state of the art, possible future challenges and research directions. An underlying primary theme of this workshop is to debate whether specialized models and evaluation approaches are necessary to properly address the challenges that arise in large scale personalization problems. The topics of interest include but are not limited to: * Is it necessary to develop fundamentally new approaches and evaluation strategies to properly address personalization? * What are appropriate objective/evaluation metrics for personalization in various domains (e.g.; ads personalization, news personalization)? * How can social network information contribute to personalization? * What breaks/what works when moving from small to large-scale personalization? * Real-time model adaptation and evaluation approaches. Online learning of personalization models. How fast can we learn personalized models? * How can learning models address the cold-start problem? * Personalization with constraints, such as budget or diversity constraints. * Privacy considerations: How much personalization is possible or acceptable?

WEBSITE:

https://sites.google.com/site/nips13personalization/

TUESDAY WORKSHOPS

TUESDAY WORKSHOP
7:30AM - 6:30PM

Acquiring and Analyzing the Activity of Large Neural Ensembles

ORGANIZERS:

Srini Turaga	Gatsby Unit, UCL
Lars Buesing	Gatsby Unit, UCL
Maneesh Sahani	Gatsby Unit, UCL
Jakob Macke	MPI Biological Cybernetics

LOCATION:

Emerald Bay 4
Harveys Convention Center Floor (CC)

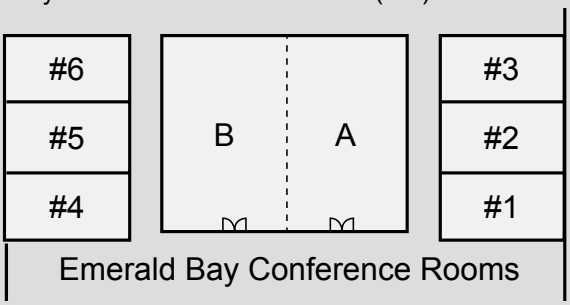

ABSTRACT:

For many years, measurements of neural activity have either been restricted to recordings from single neurons or a very small number of neurons, and anatomical reconstructions to very sparse and incomplete neural circuits. Major advances in optical imaging (e.g. 2-photon and light-sheet microscopic imaging of calcium signals) and new electrode array technologies are now beginning to provide measurements of neural activity at an unprecedented scale. High-profile initiatives such as BRAIN (Brain Research through Advancing Innovative Neurotechnologies) will fuel the development of ever more powerful techniques for mapping the structure and activity of neural circuits. Computational tools will be important to both the high-throughput acquisition of these large-scale datasets and in the analysis. Acquiring, analyzing and integrating these sources of data raises major challenges and opportunities for computational neuroscience and machine learning: i) What kind of data will be generated by large-scale functional measurements in the next decade? How will it be quantitatively or qualitatively different to the kind of data we have had previously? ii) Algorithmic methods have played an important role in data acquisition, e.g. spike-sorting algorithms or spike-inference algorithms from calcium traces. In the future, what role will computational tools play in the process of high-throughput data acquistion? iii) One of the key-challenges is to link anatomical with functional data -- what computational analysis tools will help in providing a link between these two disparate source of data? What can we learn by measuring 'functional connectivity'? iv) What have we really learned from high-dimensional recordings that is new? What will we learn? What theories could we test, if only we had access to recordings from more neurons at the same time? We have invited scientists whose research addresses these questions including prominent technologists, experimental neuroscientists, theorists and computational neuroscientists. We foresee active discussions amongst this multi-disciplinary group of scientists to catalyze exciting new research and collaborations.

Partial funding for this workshop will be provided by the Bernstein Center for Computational Neuroscience Tübingen.

WEBSITE:

http://www.bccn-tuebingen.de/events/nips-workshop-2013.html

TUESDAY WORKSHOP
7:30AM - 6:30PM

Bayesian Optimization in Theory and Practice

ORGANIZERS:

Matthew Hoffman	University of Cambridge
Jasper Snoek	Harvard University
Nando de Freitas	Oxford & UBC
Michael Osborne	University of Oxford

LOCATION:

Emerald Bay A
Harveys Convention Center Floor (CC)

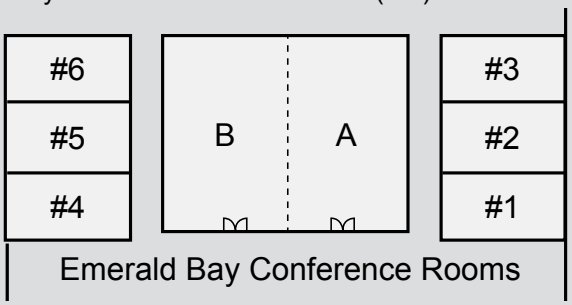

ABSTRACT:

There have been many recent advances in the development of machine learning approaches for active decision making and optimization. These advances have occurred in seemingly disparate communities, each referring to the problem using different terminology: Bayesian optimization, experimental design, bandits, active sensing, automatic algorithm configuration, personalized recommender systems, etc. Recently, significant progress has been made in improving the methodologies used to solve high-dimensional problems and applying these techniques to challenging optimization tasks with limited and noisy feedback. This progress is particularly apparent in areas that seek to automate machine learning algorithms and website analytics. Applying these approaches to increasingly harder problems has also revealed new challenges and opened up many interesting research directions both in developing theory and in practical application. Following on last year's NIPS workshop, "Bayesian Optimization & Decision Making", the goal of this workshop is to bring together researchers and practitioners from these diverse subject areas to facilitate cross-fertilization by discussing challenges, findings, and sharing data. This year we plan to focus on the intersection of "Theory and Practice". Specifically, we would like to carefully examine the types of problems where Bayesian optimization performs well and ask what theoretical guarantees can be made to explain this performance? Where is the theory lacking? What are the most pressing challenges? In what way can this empirical performance be used to guide the development of new theory?

WEBSITE:

www.bayesianoptimization.org

Constructive Machine Learning

ORGANIZERS:

Thomas Gaertner	Fraunhofer IAIS & University of Bonn
Roman Garnett	University of Bonn
Andrea Passerini	Università degli Studi di Trento

LOCATION:

Emerald Bay 1
Harveys Convention Center Floor (CC)

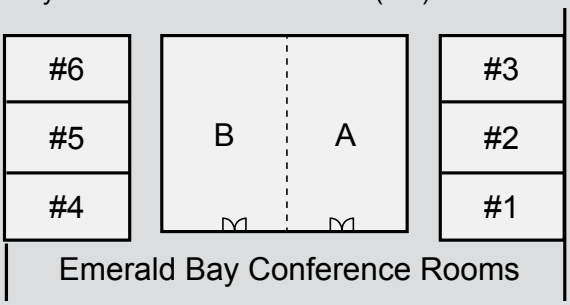

ABSTRACT:

In many real-world applications, machine learning algorithms are employed as a tool in a “constructive process”. These processes are similar to the general knowledge-discovery process but have a more specific goal: the construction of one-or-more domain elements with particular properties. The most common use of machine learning algorithms in this context is to predict the properties of candidate domain elements. In this workshop we want to bring together domain experts employing machine learning tools in constructive processes and machine learners investigating novel approaches or theories concerning constructive processes as a whole. The concerned machine learning approaches are typically interactive (e.g., online- or active-learning algorithms) and have to deal with huge, relational in- and/or output spaces. Interesting applications include but are not limited to: de novo drug design, generation of art (e.g., music composition), construction of game levels, generation of novel food recipes, proposal of travel itineraries, etc. Interesting approaches include but are not limited to: active approaches to structured output learning, transfer or multi-task learning of generative models, active search or online optimisation over relational domains, and learning with constraints. Many of the applications of constructive machine learning, including the ones mentioned above, are primarily considered in their respective application domain research area but are hardly present at machine learning conferences. By bringing together domain experts and machine learners working on constructive ML, we hope to bridge this gap between the communities.

WEBSITE:

http://www-kd.iai.uni-bonn.de/cml/

TUESDAY WORKSHOP
7:30AM - 6:30PM

Data Driven Education

ORGANIZERS:

Jonathan Huang	Stanford University
Sumit Basu	Microsoft Research
Kalyan Veeramachaneni	MIT

LOCATION:

Tahoe D
Harrah's Special Events Center 2nd Floor

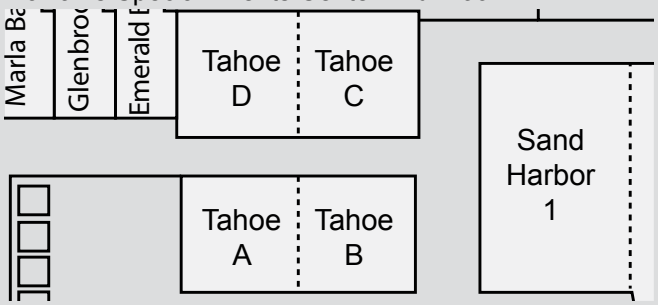

ABSTRACT:

Given the incredible technological leaps that have changed so many aspects of our lives in the last hundred years, it's surprising that our approach to education today is much the same as it was a century ago. While successful educational technologies have been developed and deployed in some areas, we have yet to see a widespread disruption in teaching methods at the primary, secondary, or post-secondary levels. However, as more and more people gain access to broadband internet, and new technology-based learning opportunities are introduced, we may be witnessing the beginnings of a revolution in educational methods. With college tuitions rising, school funding dropping, test scores falling, and a steadily increasing world population desiring high-quality education at low cost, the impact of educational technology seems more important than ever. With these technology-based learning opportunities, the rate at which educational data is being collected has also exploded in recent years as an increasing number of students have turned to online resources, both at traditional universities as well as massively open-access online courses (MOOCs) for formal or informal learning. This change raises exciting challenges and possibilities particularly for the machine learning and data sciences communities. These trends and changes are the inspiration for this workshop, and our first goal is to highlight some of the exciting and impactful ways that our community can bring tools from machine learning to bear on educational technology. Some examples include (but are not limited to) the following: + Adaptive and personalized education + Assessment: automated, semi-automated, and peer grading + Gamification and crowdsourcing in learning + Large scale analytics of MOOC data + Multimodal sensing + Optimization of pedagogical strategies and curriculum design + Content recommendation for learners + Interactive Tutoring Systems + Intervention evaluations and causality modeling + Supporting collaborative and social learning + Data-driven models of human learning The second goal of the workshop is to accelerate the progress of research in these areas by addressing the challenges of data availability. At the moment, there are several barriers to entry including the lack of open and accessible datasets as well as unstandardized formats for such datasets. We hope that by (1) surveying a number of the publicly available datasets, and (2) proposing ways to distribute other datasets such as MOOC data in a spirited panel discussion we can make real progress on this issue as a community, thus lowering the barrier for researchers aspiring to make a big impact in this important area. Target Audience: + Researchers interested in analyzing and modeling educational data, + Researchers interested in improving or developing new data-driven educational technologies, + Others from the NIPS community curious about the trends in online education and the opportunities for machine learning research in this rapidly-developing area.

WEBSITE:

http://lytics.stanford.edu/datadriveneducation

TUESDAY WORKSHOP
7:30AM - 6:30PM

Greedy Algorithms, Frank-Wolfe and Friends: A Modern Perspective

ORGANIZERS:

Martin Jaggi	UC Berkeley
Zaid Harchaoui	INRIA Grenoble, France
Federico Pierucci	INRIA Grenoble

LOCATION:

Emerald Bay 6
Harveys Convention Center Floor (CC)

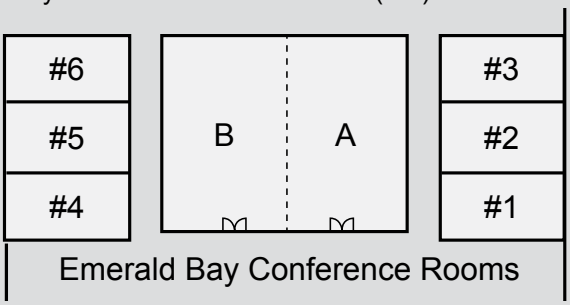

ABSTRACT:

Greedy algorithms and projection-free first-order optimization algorithms are at the core of many of the state of the art sparse methods in machine learning, signal processing, harmonic analysis, statistics and other seemingly unrelated areas, with different goals at first sight. Examples include matching pursuit, boosting, greedy methods for sub-modular optimization, with applications ranging from large-scale structured prediction to recommender systems. In the field of optimization, the recent renewed interest in Frank-Wolfe/conditional gradient algorithms opens up an interesting perspective towards a unified understanding of these methods, with a big potential to translate the rich existing knowledge about the respective greedy methods between the different fields. The goal of this workshop is to take a step towards building a modern and consistent perspective on these related algorithms. The workshop will gather renowned experts working on those algorithms in machine learning, optimization, signal processing, statistics and harmonic analysis, in order to engender a fruitful exchange of ideas and discussions and to push further the boundaries of scalable and efficient optimization for learning problems.

WEBSITE:

http://sites.google.com/site/nips13greedyfrankwolfe

Knowledge Extraction from Text (KET)

ORGANIZERS:

Marko Grobelnik	Jozef Stefan Institute
Blaz Fortuna	Jozef Stefan Institute
Estevam Hruschka	UFSCar
Michael Witbrock	Cycorp Inc

LOCATION:

Emerald Bay 2
Harveys Convention Center Floor (CC)

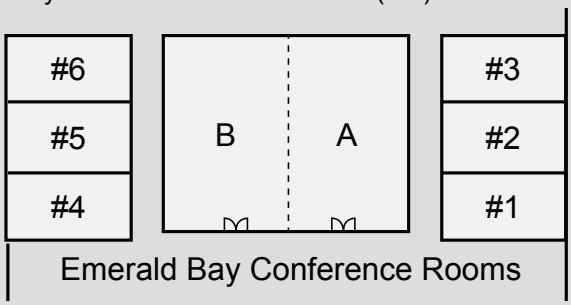

ABSTRACT:

Text understanding is an old yet-unsolved AI problem consisting of a number of nontrivial steps. The critical step in solving the problem is knowledge acquisition from text, i.e. a transition from a non-formalized text into a formalized actionable language (i.e. capable of reasoning). Other steps in the text understanding pipeline include linguistic processing, reasoning, text generation, search, question answering etc. which are more or less solved to the degree which allows composition of a text understanding service. On the other hand, we know that knowledge acquisition, as the key bottleneck, can be done by humans, while automating of the process is still out of reach in its full breadth. After failed attempts in the past (due to a lack of theoretical and technological prerequisites), in the recent years the interest for the text understanding and knowledge acquisition form text is growing. There is a number of AI research groups dealing with the various aspects in the areas of computational linguistics, machine learning, probabilistic & logical reasoning, and semantic web. The commonality among all the newer approaches is the use of machine learning to deal with representational change. To list some of the groups working in the area:

- Carnegie Mellon University (Never-Ending Language Learning: http://rtw.ml.cmu.edu/rtw/)
- Cycorp (Semantic Construction Grammar: http://www.cyc.com/)
- IBM Research (Watson project: http://www.ibm.com/watson)
- IDIAP Research Institute (Deep Learning for NLP: http://publications.idiap.ch/index.php/authors/show/336)
- Jozef Stefan Institute (Cross-Lingual Knowledge-Extraction: http://xlike.org)
- KU Leuven (Spatial Role Labelling via Machine Learning for SEMEVAL)
- Max Planck Institut (YAGO project: http://www.mpi-inf.mpg.de/yago-naga/yago/)
- MIT Media Lab (ConceptNet: http://conceptnet5.media.mit.edu/)
- University Washington (Open Information Extraction: http://openie.cs.washington.edu/)
- Vulcan Inc. (Semantic Inferencing on Large Knowledge: http://silk.semwebcentral.org/)

Apart from the above projects, there is noticeable increase of interest in the technology companies (such as Google, Microsoft, IBM) as well as big publishers (such as NYTimes, BBC, Bloomberg) to employ semantic technologies into their services leading towards understanding unstructured data beyond shallow, representation poor Text-Mining and Information-Retrieval techniques. Workshop objective: Since all of the above listed attempts use extensively machine learning and probabilistic approaches, the goal of the workshop is to collect key researchers and practitioners from the area to exchange ideas, approaches and techniques used to deal with text understanding and related knowledge acquisition problems.

WEBSITE:

Learning Faster From Easy Data

ORGANIZERS:

Peter Grunwald	CWI
Wouter Koolen	Queensland University of Technology
Sasha Rakhlin	University of Pennsylvania

LOCATION:

Tahoe A
Harrah's Special Events Center 2nd Floor

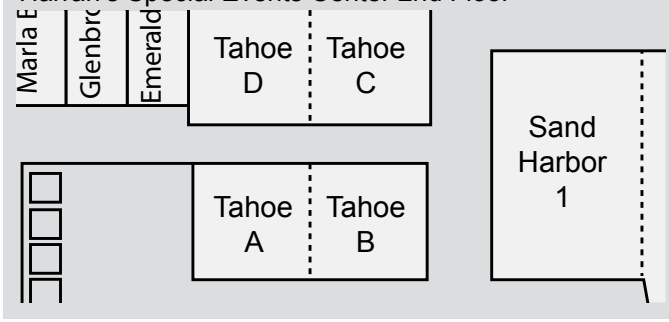

ABSTRACT:

Most existing theory in both online and statistical learning is centered around a worst-case analysis. For instance, in online learning data are assumed to be generated by an adversary and the goal is to minimize regret. In statistical learning the majority of theoretical results consider risk bounds for the worst-case i.i.d. data generating distribution. In both cases the worst case convergence rates (for regret/n and risk) for 0/1-type and absolute loss functions are O(1/sqrt{n}). Yet in practice simple heuristics like Follow-the-Leader (FTL) often empirically exhibit faster rates. It has long been known that under Vovk's (1990) mixability condition on the loss function, faster rates are possible. Even without mixability or the closely related exp-concavity (Cesa-Bianchi and Lugosi 2006), in the statistical setting there exist conditions on the distribution under which faster learning rates can be obtained; the main example being Tsybakov's (2004) margin condition, which was recently shown to be intimately connected to mixability (Van Erven et al., 2012). In practice, even if the loss is not mixable and no distributional assumptions apply, the data are nevertheless often easy enough to allow accelerated learning. Initial promising steps in this direction have been made recently, including parameterless algorithms that combine worst-case O(1/sqrt{n}) regret guarantees for the adversarial setting with - fast rates in the stochastic bandit setting (Bubeck and Slivkins, COLT 2012) - exploitation of observably sub-adversarial data (Rakhlin, Shamir and Sridharan, AISTATS 2013) - learning as fast as FTL whenever FTL works well (De Rooij, Van Erven, Grünwald and Koolen, JMLR 2013) It remains a huge challenge however, to characterize the types of data for which faster learning is possible, to define `easy data' in a generic way, let alone to design algorithms that automatically adapt to exploit it. The aim of this day-long workshop is threefold 1) to map, by means of a series of invited and contributed talks, the existing landscape of "easiness criteria" in relation to the efficiency of their corresponding algorithms, 2) to identify, by means of a panel discussion led by the organizers, obstacles and promising directions, 3) and through interaction foster partnerships for future research. Discussion will be centered around the so-far elusive concept of easy data. Can the existing characterizations based on variances, mixability gaps, FTL etc. be brought under a common umbrella? Can ideas and approaches from statistical learning theory be transported to online learning (and vice versa)?

WEBSITE:

http://www.cwi.nl/~wmkoolen/easydata2013/

TUESDAY WORKSHOP
7:30AM - 6:30PM

Machine Learning for Clinical Data Analysis and Healthcare

ORGANIZERS:

Jenna Wiens	MIT
Finale Doshi-Velez	Harvard
Can Ye	Carnegie Mellon University
Madalina Fiterau	Carnegie Mellon University
Shipeng Yu	Siemens Healthcare
Le Lu	National Institutes of Health
Balaji Krishnapuram	Siemens Medical Solutions USA, Inc.

LOCATION:

Tallac
Harveys 12th Floor “Top Of The Wheel”

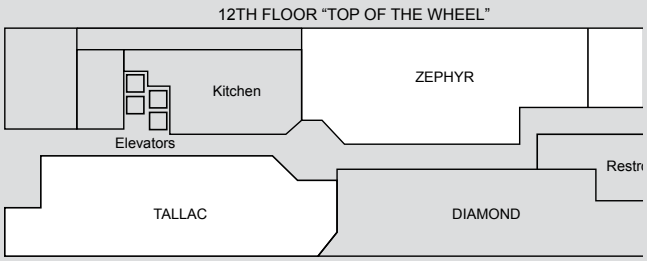

ABSTRACT:

Advances in medical information technology have resulted in enormous warehouses of data that are at once overwhelming and sparse. A single patient visit may result in tens to thousands of measurements and structured information, including clinical factors, diagnostic imaging, lab tests, genomic and proteomic tests. Hospitals may see thousands of patients each year. However, each patient may have relatively few visits to any particular medical provider. The resulting data are a heterogeneous amalgam of patient demographics, vital signs, diagnoses, records of treatment and medication receipt and annotations made by nurses or doctors, each with its own idiosyncrasies. The objective of this workshop is to discuss how advanced machine learning techniques can derive clinical and scientific impact from these messy, incomplete, and partial data. We will bring together machine learning researchers and experts in medical informatics who are involved in the development of algorithms or intelligent systems designed to improve quality of healthcare. Relevant areas include health monitoring systems, clinical data labeling and clustering, clinical outcome prediction, efficient and scalable processing of medical records, feature selection or dimensionality reduction in clinical data, tools for personalized medicine, and time-series analysis with medical applications.

WEBSITE:

http://goo.gl/lpgjyN

Machine Learning for Sustainability

ORGANIZERS:

Edwin Bonilla	NICTA
Tom Dietterich	Oregon State University
Theo Damoulas	New York University
Andreas Krause	ETH Zurich

LOCATION:

Glenbrook+Emerald
Harrah's Special Events Center 2nd Floor

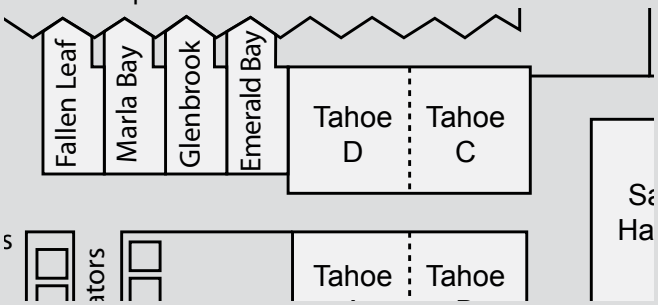

ABSTRACT:

Sustainability encompasses the balance of environmental, economic and societal demands. There is strong evidence suggesting that more actions need to be taken in order to achieve this balance. For example, Edward O. Wilson said in his 2002 Book The Future of Life that "at the current rates of human destruction of natural ecosystems, 50% of all species of life on earth will be extinct in 100 years". More recently, a 2012 review in Nature has stated that, similarly to localized ecological systems, "the global ecosystem as a whole can react in the same way and is approaching a planetary-scale critical transition as a result of human influence". While the significance of the problem is apparent, more involvement from the machine learning community in sustainability problems is required. Not surprisingly, sustainability problems bring along interesting challenges and opportunities for machine learning in terms of complexity, scalability and impact in areas such as prediction, modeling and control. This workshop aims at bringing together scientists in machine learning, operations research, applied mathematics and statistics with a strong interest in sustainability to discuss how to use existing techniques and how to develop novel methods in order to address such challenges. There are many application areas in sustainability where machine learning can have a significant impact. For example: - Climate change - Conservation and biodiversity - Socio-economic systems - Understanding energy consumption - Renewable energy - Impact of mining - Sustainability in the developing world - Managing the power grid - Biofuels Similarly, machine learning approaches to sustainability problems can be drawn from several fields such as: - Constraint optimization - Dynamical systems - Spatio-temporal modeling - Probabilistic inference - Sensing and monitoring - Decision making under uncertainty - Stochastic optimization The talks and posters are expected to span (but not be limited to) the above areas. More importantly, there will be a specific focus on how cutting-edge machine learning research is developed (i.e. not only using off-the-shelf ML techniques) in order to address challenges in terms of complexity, scalability and impact that sustainability problems may pose. The main expected outcomes of this workshop are: (1) attracting more people to work on computational sustainability; (2) transfer of knowledge across different application domains; and (3) emerging collaboration between participants. More long-term avenues such as datasets and competitions will be explored. There will be an award (~ $250 book voucher) for the best contribution, which will be given an oral presentation.

WEBSITE:

https://sites.google.com/site/mlsustws/

TUESDAY WORKSHOP
7:30AM - 6:30PM

Machine Learning in Computational Biology

ORGANIZERS:

Jean-Philippe Vert	ParisTech
Anna Goldenberg	University of Toronto
Sara Mostafavi	Stanford University
Oliver Stegle	EMBL-EBI

LOCATION:

Zephyr
Harveys 12th Floor “Top Of The Wheel”

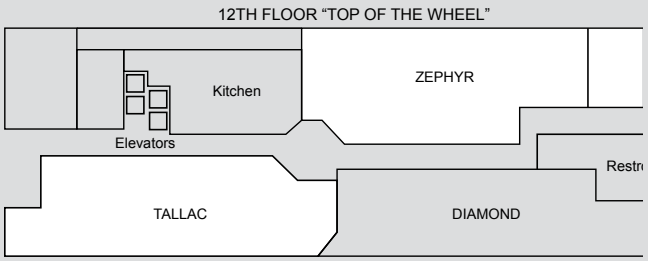

ABSTRACT:

The field of computational biology has seen dramatic growth over the past few years, both in terms of new available data, new scientific questions, and new challenges for learning and inference. In particular, biological data are often relationally structured and highly diverse, well-suited to approaches that combine multiple weak evidence from heterogeneous sources. These data may include sequenced genomes of a variety of organisms, gene expression data from multiple technologies, protein expression data, protein sequence and 3D structural data, protein interactions, gene ontology and pathway databases, genetic variation data (such as SNPs), and an enormous amount of textual data in the biological and medical literature. Furthermore, next generation sequencing technologies and high-throughput imaging techniques are yielding terabyte scale data sets that require novel algorithmic solutions. New types of scientific and clinical problems require the development of novel supervised and unsupervised learning methods that can use these growing resources. The goal of this workshop is to present emerging problems and innovative machine learning techniques in computational biology. We will invite several speakers from the biology/bioinformatics community who will present current research problems in computational biology, and we will invite contributed talks on novel learning approaches in computational biology. We encourage contributions describing either progress on new bioinformatics problems or work on established problems using methods that are substantially different from established alternatives. Kernel methods, graphical models, feature selection, non-parametric models and other techniques applied to relevant bioinformatics problems would all be appropriate for the workshop. We are particularly keen on considering contributions related to the prediction of functions from genotypes and to applications in personalized medicine, as illustrated by our invited speakers. The targeted audience are people with interest in learning and applications to relevant problems from the life sciences, including NIPS participants without any existing research link to computational biology.

WEBSITE:

http://www.mlcb.org

Machine Learning Open Source Software: Towards Open Workflows

ORGANIZERS:

Antti Honkela University of Helsinki
Cheng Soon Ong NICTA, Melbourne

LOCATION:

Emerald Bay 5
Harveys Convention Center Floor (CC)

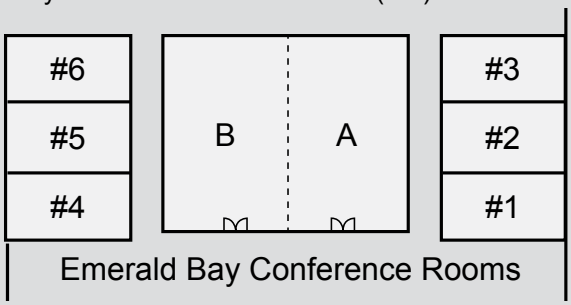

ABSTRACT:

Machine learning open source software (MLOSS) is one of the cornerstones of open science and reproducible research. Along with open access and open data, it enables free reuse and extension of current developments in machine learning. The mloss.org site exists to support a community creating a comprehensive open source machine learning environment, mainly by promoting new software implementations. This workshop aims to enhance the environment by fostering collaboration with the goal of creating tools that work with one another. Far from requiring integration into a single package, we believe that this kind of interoperability can also be achieved in a collaborative manner, which is especially suited to open source software development practices. The workshop is aimed at all machine learning researchers who wish to have their algorithms and implementations included as a part of the greater open source machine learning environment. Continuing the tradition of well received workshops on MLOSS at NIPS 2006, NIPS 2008 and ICML 2010, we will have a workshop that is a mix of invited speakers, contributed talks and demos as well as a discussion session. For 2013, we focus on workflows and pipelines. Many algorithms and tools have reached a level of maturity which allows them to be reused and integrated into larger systems.

WEBSITE:

http://mloss.org/workshop/nips13/

TUESDAY WORKSHOP
7:30AM - 6:30PM

MLINI-13: Machine Learning And Interpretation In Neuroimaging (Day 2)

ORGANIZERS:

Irina Rish	IBM T.J. Watson Research Ctr
Georg Langs	Medical University of Vienna
Brian Murphy	Carnegie Mellon University
Kai-min Chang	Carnegie Mellon University
Guillermo Cecchi	IBM Research

LOCATION:

Sierra
Harveys Casino Main Floor

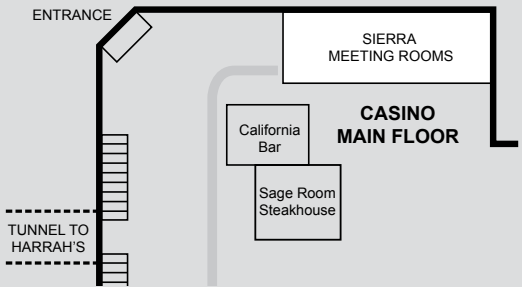

ABSTRACT:

This is a two-day workshop on the topic of machine learning approaches in neuroscience, neuroimaging, with a specific extension to behavioral experiments and psychology. We believe that machine learning has a prominent role in shaping how questions in neuroscience are framed, and that the machine-learning mind set is now entering modern psychology and behavioral studies. The goal of the workshop is to pinpoint the most pressing issues and common challenges across the neuroscience and machine learning, and to sketch future directions and open questions in the light of novel methodology. The workshop is aimed at offering a forum that joins machine learning, neuroscience, and psychology community, and should facilitate formulating and discussing the issues at their interface.

Motivated by two previous workshops, MLINI '11 and MLINI'12, we will center this workshop around invited talks, and two panel discussions. Triggered by these discussions, this year we plan to adapt the workshop topics to a less traditional scope that investigates the role of machine learning in neuroimaging of both animals and humans, as well as in behavioral models and psychology.

We intend to focus on the implications that follow from adopting multivariate machine-learning methods for studying brain function. In particular, we would like to investigate how these methods may be used to represent cognitive states, and what ramifications this has for consequent theories of cognition. Besides providing a rationale for the use of machine-learning methods in studying brain function, a further goal of this workshop is to identify shortcomings of state-of-the-art approaches and initiate research efforts that increase the impact of machine learning on cognitive neuroscience. Possible topics for discussion include the following main themes: I) machine learning and pattern recognition methodology in brain research, II) functional connectivity and dynamical models of brain activity, III) multi-modal analysis including mental state inference from behavioral data, and IV) linking machine learning, neuroimaging and neuroscience

WEBSITE:

https://sites.google.com/site/mlininips2013/

TUESDAY WORKSHOP
7:30AM - 6:30PM

Neural Information Processing Scaled for Bioacoustics

ORGANIZERS:

Hervé Glotin	Toulon University & Institut Univ. France
Yann LeCun	New York University
Thierry Artières	LIP6/UPMC
Stephane Mallat	LIENS, Ecole Normale Supérieure
Ofer Tchernichovski	Hunter College, CUNY
Xanadu Halkias	Université de Toulon

LOCATION:

Tahoe C
Harrah's Special Events Center 2nd Floor

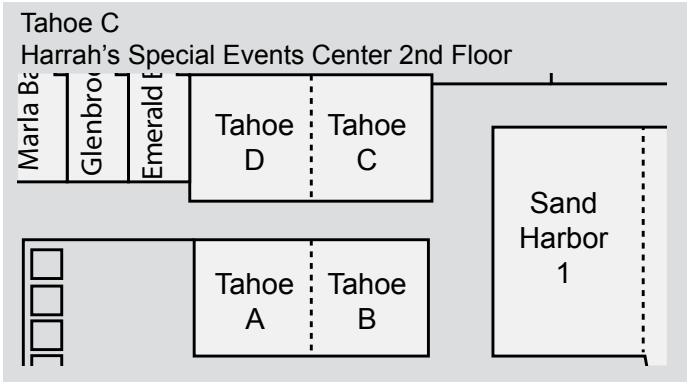

ABSTRACT:

Bioacoustic data science aims at modeling animal sounds for neuroethology and biodiversity assessment. It has received increasing attention due to its diverse potential benefits. It is steadily required by regulatory agencies for timely monitoring of environmental impacts from human activities. Given the complexity of the collected data along with the numerous species and environmental contexts, bioacoustics requires robust information processing.

The features and biological significance of animal sounds, are constrained by the physics of sound production and propagation, and evolved through the processes of natural selection. This yields to new paradigms such as curriculum song learning, predator-prey acoustic loop, etc. NIPS4B solidifies an innovative computational framework by focusing on the principles of information processing, if possible in an inheretly hierarchical manner or with physiological parallels: Deep Belief Networks (DBN), Sparse Auto Encoders (SAE), Convolutional Networks (ConvNet), Scattering transforms etc. It encourages interdisciplinary, scientific exchanges and foster collaborations, bringing together experts from machine learning and computational auditory scene analysis, within animal sound and communication systems.

One challenge concerns bird classification (on Kaggle): identify 87 species of Provence (recordings Biotope SA). It is the biggest bird song challenge according to our knowledge, more complex than ICML4B (sabiod.org/ICML4B2013_proceedings.pdf). A second challenge concerns the representation of a remarkable humpback whale song (Darewin - La Reunion), in order to help its analysis. Other special session concerns (neural) modelisation of the biosonar of bats or dolphins.

WEBSITE:

http://sabiod.org/nips4b

References:

Glotin H, Dugan P, LeCun Y, Clark C, Halkias X, (2013) Proc. of the first workshop on Machine Learning for Bioacoustics, sabiod.org/ICML4B2013_proceedings.pdf, ICML4B

Glotin H, (2013) Etho-Acoustics: Categorisation & Localisation into Soundscapes, Ed. Intech open book

Pace F, Benard F, Glotin H, Adam O, White P, (2010) Subunit definition for humpback whale call classification, J. Applied Acoustics, 11(71)

Halkias X, Paris S, Glotin H, (2013) Classification of mysticete sounds using machine learning techniques, 134, 3496, JASA

Lallemand I, Schwarz D, Artieres T, (2012) A Multiresolution Kernel Distance for SVM Classification of Environmental Sounds, SMC

Farabet C, Couprie C, Najman L, LeCun Y, (2013) Learning Hierarchical Features for Scene Labeling, IEEE PAMI

LeCun, Y, Learning Invariant Feature Hierarchies, (2012) Workshop on Biological & Computer Vision Interfaces, LNCS, V7583, ECCV

Anden J, Mallat S, (2011) Scattering transform applied to audio signals & musical classification: Multiscale Scattering for Audio Classification, ISMIR

Lipkind D, Marcus GF...Tchernichovski O, (2013) Stepwise acquisition of vocal combinatorial capacity in songbirds & human infants, 10.1038/nature12173, Nature

Tchernichovski O, Wallman J, (2008) Neurons of imitation, 451(17), Nature

Lallemand I, Schwarz D, Artieres T, (2012) A Multiresolution Kernel Distance for SVM Classification of Environmental Sounds, SMC

Soullard Y, Artieres T, (2011) Hybrid HMM and HCRF model for sequence classification, ESANN

TUESDAY WORKSHOP
7:30AM - 6:30PM

New Directions in Transfer and Multi-Task: Learning Across Domains and Tasks

ORGANIZERS:

Urun Dogan	Microsoft
Marius Kloft	Courant Institute, NYU & MSKCC
Tatiana Tommasi	KUL
Francesco Orabona	Toyota Technological Institute Chicago

LOCATION:

Fallen Leaf + Marla Bay
Harrah's Special Events Center 2nd Floor

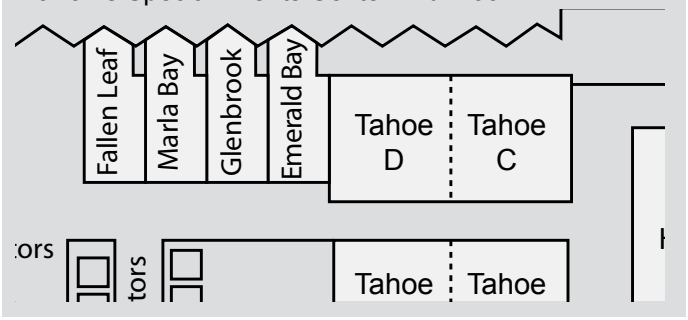

ABSTRACT:

The main objective of the workshop is to document and discuss the recent rise of new research questions on the general problem of learning across domains and tasks. This includes the main topics of transfer [1,2,3] and multi-task learning [4], together with several related variants as domain adaptation [5,6] and dataset bias [7]. In the last years there has been an increasing boost of activity in these areas, many of them driven by practical applications, such as object categorization. Different solutions were studied for the considered topics, mainly separately and without a joint theoretical framework. On the other hand, most of the existing theoretical formulations model regimes that are rarely used in practice (e.g. adaptive methods that store all the source samples). The workshop will focus on closing this gap by providing an opportunity for theoreticians and practitioners to get together in one place, to share and debate over current theories and empirical results. The goal is to promote a fruitful exchange of ideas and methods between the different communities, leading to a global advancement of the field. Transfer Learning - Transfer Learning (TL) refers to the problem of retaining and applying the knowledge available for one or more source tasks, to efficiently develop an hypothesis for a new target task. Each task may contain the same (domain adaptation) or different label sets (across category transfer). Most of the effort has been devoted to binary classification, while most interesting practical transfer problems are intrinsically multi-class and the number of classes can often increase in time. Hence, it is natural to ask: - How to formalize knowledge transfer across multi-class tasks and provide theoretical guarantees on this setting? - Moreover, can interclass transfer and incremental class learning be properly integrated? - Can learning guarantees be provided when the adaptation relies only on pre-trained source hypotheses without explicit access to the source samples, as it is often the case in real world scenarios? Multi-task Learning - Learning over multiple related tasks can outperform learning each task in isolation. This is the principal assertion of Multi-task learning (MTL) and implies that the learning process may benefit from common information shared across the tasks. In the simplest case, transfer process is symmetric and all the tasks are considered as equally related and appropriate for joint training. - What happens when this condition does not hold, e.g., how to avoid negative transfer? - Moreover, can RHKS embeddings be adequately integrated into the learning process to estimate and compare the distributions underlying the multiple tasks? - How may embedding probability distributions help learning from data clouds? - Recent methods, like deep learning or multiple kernel learning, can help to get a step closer towards the complete automatization of multi-task learning? - How can notions from reinforcement learning such as source task selection be connected to notions from convex multi-task learning such as the task similarity matrix?

WEBSITE:

https://sites.google.com/site/learningacross/

Resource-Efficient Machine Learning

ORGANIZERS:

Yevgeny Seldin	Queensland Univ. & UC Berkeley
Yasin Abbasi	Queensland Univ. of Technology
Koby Crammer	he Technion
Ralf Herbrich	Amazon Development Center
Peter Bartlett	UC Berkeley

LOCATION:

Emerald Bay 3
Harveys Convention Center Floor (CC)

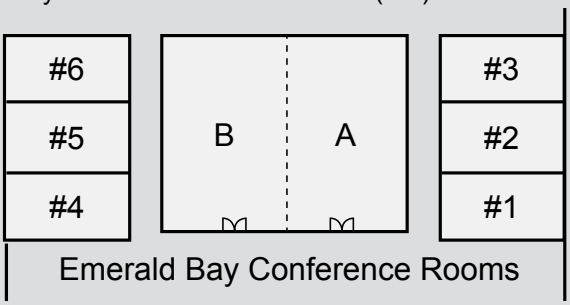

ABSTRACT:

Resource efficiency is crucial in many tasks, ranging from large-scale learning (“big data”) to small-scale mobile devices. Understanding resource efficiency is also important for understanding biological systems, from individual cells to complex learning systems, such as the human brain. The goal of this workshop is to improve our fundamental theoretical understanding and link between various applications of learning under constraints on the resources, such as computation, observations, communication, and memory. While the founding fathers of machine learning were mainly concerned with characterizing the sample complexity of learning (the observations resource) it now gets realized that fundamental understanding of other resource requirements, such as computation, communication, and memory is equally important for further progress.

We believe that there are deep connections between problems at various scales and with various resource constraints and that there are basic principles of learning under resource constraints and new optimization algorithms that are yet to be discovered. We invite researchers to share their practical challenges and theoretical insights into this problem.

WEBSITE:

https://sites.google.com/site/resefml2013

TUESDAY WORKSHOP
7:30AM - 6:30PM

Topic Models: Computation, Application And Evaluation

ORGANIZERS:

David Mimno	Cornell University
Amr Ahmed	Research at Google
David Blei	Princeton University
Jordan Boyd-Graber	University of Maryland
Ankur Moitra	IAS
Hanna Wallach	Univ. of Massachusetts Amherst

LOCATION:

Emerald Bay B
Harveys Convention Center Floor (CC)

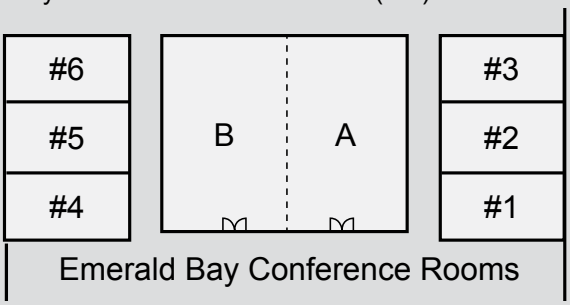

ABSTRACT:

Since the most recent NIPS topic model workshop in 2010, interest in statistical topic modeling has continued to grow in a wide range of research areas, from theoretical computer science to English literature. The goal of this workshop, which marks the 10th anniversary of the original LDA NIPS paper, is to bring together researchers from the NIPS community and beyond to share results, ideas, and perspectives. We will organize the workshop around the following three themes: Computation: The computationally intensive process of training topic models has been a useful testbed for novel inference methods in machine learning, such as stochastic variational inference and spectral inference. Theoretical computer scientists have used LDA as a test case to begin to establish provable bounds in unsupervised machine learning. This workshop will provide a forum for researchers developing new inference methods and theoretical analyses to present work in progress, as well as for practitioners to learn about state of the art research in efficient and provable computing. Applications: Topic models are now commonly used in a broad array of applications to solve real-world problems, from questions in digital humanities and computational social science to e-commerce and government science policy. This workshop will share new application areas, and discuss our experiences adapting general tools to the particular needs of different settings. Participants will look for commonalities between diverse applications, while also using the particular challenges of each application to define theoretical research agendas. Evaluation: A key strength of topic modeling is its exceptional capability for exploratory analysis, but evaluating such use can be challenging: there may be no single right answer. As topic models become widely used outside machine learning, it becomes increasingly important to find evaluation strategies that match user needs. The workshop will focus both on the specifics of individual evaluation metrics and the more general process of iteratively criticizing and improving models. We will also consider questions of interface design, visualization, and user experience.

Program committee (confirmed): Edo Airoldi (Harvard), Laura Dietz (UMass), Jacob Eisenstein (GTech), Justin Grimmer (Stanford), Yoni Halpern (NYU), Daniel Hsu (Columbia), Brendan O'Connor (CMU), Michael Paul (JHU), Eric Ringger (BYU), Brandon Stewart (Harvard), Chong Wang (CMU), Sinead Williamson (UT-Austin)

WEBSITE:

http://nips2013.topicmodels.net

Workshop on Spectral Learning

ORGANIZERS:

Byron Boots	University of Washington
Daniel Hsu	Columbia University
Borja Balle	McGill University

LOCATION:

Tahoe B
Harrah's Special Events Center 2nd Floor

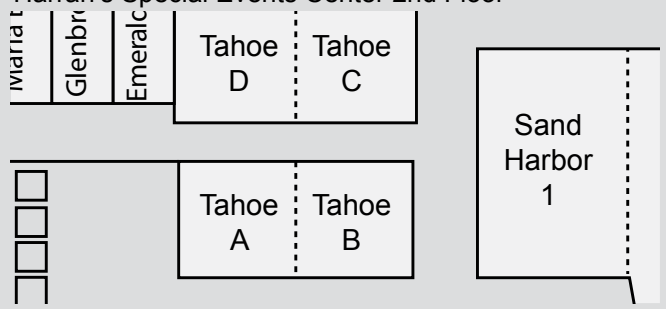

ABSTRACT:

Many problems in machine learning involve collecting high-dimensional multivariate observations or sequences of observations, and then fitting a compact model which explains these observations. Recently, linear algebra techniques have given a fundamentally different perspective on how to fit and perform inference in these models. Exploiting the underlying spectral properties of the model parameters has led to fast, provably consistent methods for parameter learning that stand in contrast to previous approaches, such as Expectation Maximization, which suffer from bad local optima and slow convergence. In the past several years, these Spectral Learning algorithms have become increasingly popular. They have been applied to learn the structure and parameters of many models including predictive state representations, finite state transducers, hidden Markov models, latent trees, latent junction trees, probabilistic context free grammars, and mixture/admixture models. Spectral learning algorithms have also been applied to a wide range of application domains including system identification, video modeling, speech modeling, robotics, and natural language processing. The focus of this workshop will be on spectral learning algorithms, broadly construed as any method that fits a model by way of a spectral decomposition of moments of (features of) observations. We would like the workshop to be as inclusive as possible and encourage paper submissions and participation from a wide range of research related to this focus. We will encourage submissions on the following themes: - How can spectral techniques help us develop fast and local minima free solutions to real world problems where existing methods such as Expectation Maximization are unsatisfactory? - How do spectral/moment methods compare to maximum-likelihood estimators and Bayesian methods, especially in terms of robustness, statistical efficiency, and computational efficiency? - What notions of spectral decompositions are appropriate for latent variable models and structured prediction problems? - How can spectral methods take advantage of multi-core/multi-node computing environments? - What computational problems, besides parameter estimation, can benefit from spectral decompositions and operator parameterizations? (For example, applications to parsing.)

WEBSITE:

https://sites.google.com/site/spectrallearningworkshop/

LAKE TAHOE DESTINATION MAP

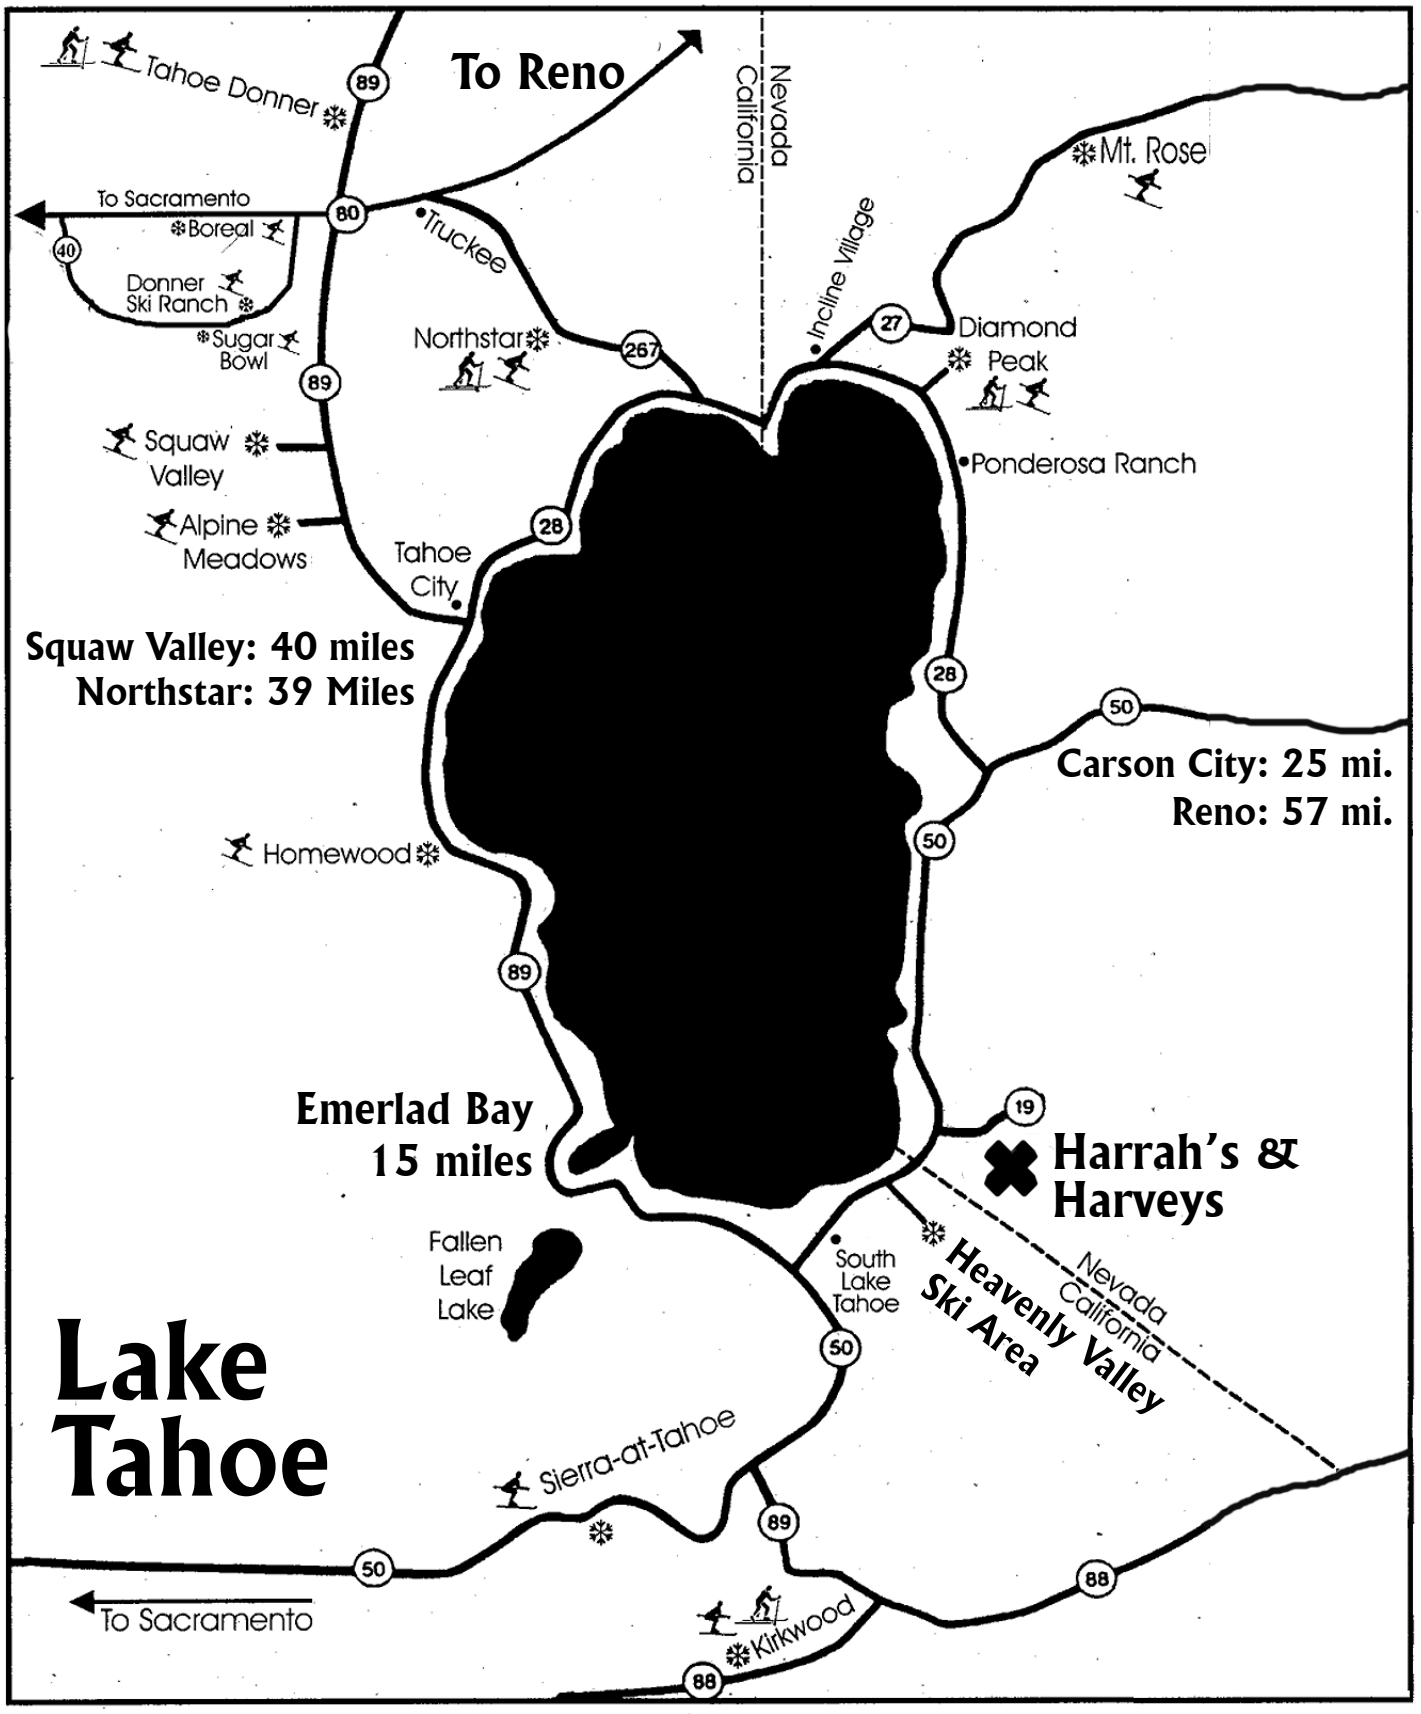

2014 NIPS Conference
Montreal • Canada
Palais des congrès de Montréal