MONTREAL 2014

WORKSHOP BOOK

2014 NIPS Conference

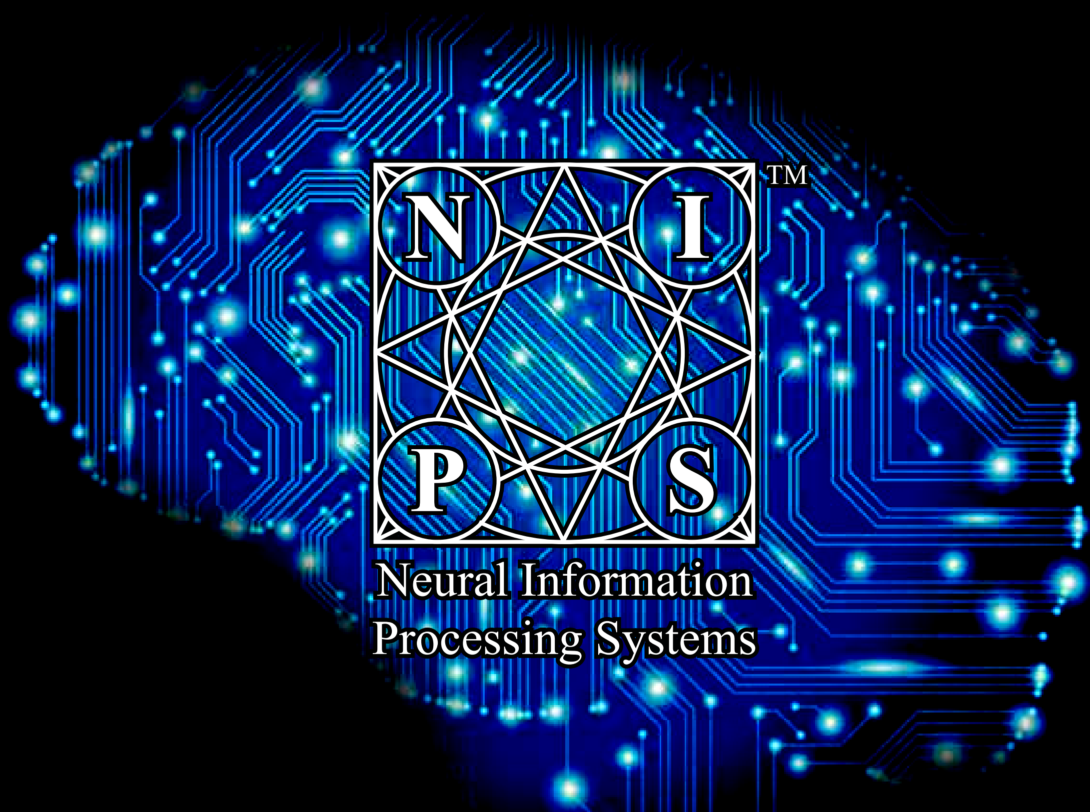

NIPS 2014

NEURAL INFORMATION PROCESSING SYSTEMS WORKSHOP

TUTORIALS
December 8, 2014
Palais des Congrès de Montréal
Convention and Exhibition Center,
Montreal, Quebec, Canada

CONFERENCE SESSIONS
December 9 - 11, 2014
Palais des Congrès de Montréal
Convention and Exhibition Center,
Montreal, Quebec, Canada

WORKSHOPS
December 12-13, 2014
Palais des Congrès de Montréal
Convention and Exhibition Center,
Montreal, Quebec, Canada

Sponsored by the Neural Information Processing System Foundation, Inc

The technical program includes 6 invited talks and 414 accepted papers, selected from a total of 1678 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 27," edited by Corinna Cortes, Neil Lawrence, Max Welling, Zoubin Ghahramani and Kilian Weinberger

Neural Information Processing Systems Foundation

TABLE OF CONTENTS

MONDAY WORKSHOPS

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

TUESDAY WORKSHOPS

2014 EXHIBITORS

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

ORGANIZING COMMITTEE

General Chairs: **Zoubin Ghahramani** (University of Cambridge)
Max Welling (University of Amsterdam)
Program Chairs: **Corinna Cortes** (Google Research)
Neil Lawrence (University of Sheffield)
Tutorials Chair: **Jeff A. Blimes** (University of Washington)
Demonstration Chair: **Hugo Larochelle** (Univ. de Sherbrooke)

Workshop Chairs: **Francis Bach** ((INRIA & ENS),
Amir Globerson (Hebrew University)
Publications Chair and Electronic
Proceedings Chair: **Kilian Weinberger** (Washington Univ)
Program Manager: **Alan Saul** (University of Sheffield)

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

EXECUTIVE BOARD
Peter Bartlett, Queensland University and UC Berkeley
Léon Bottou, Microsoft Research
Chris J.C. Burges, Microsoft Research
Fernando Pereira, Google Research
John Shawe-Taylor, University College London
Rich Zemel, University of Toronto
Chris Williams, University of Edinburgh

EMERITUS MEMBERS
Gary Blasdel, Harvard Medical School
T. L. Fine, Cornell University
Eve Marder, Brandeis University

ADVISORY BOARD
Sue Becker, McMaster University, Ontario, Canada
Yoshua Bengio, University of Montreal, Canada
Jack Cowan, University of Chicago
Thomas G. Dietterich, Oregon State University
Stephen Hanson, Rutgers University
Michael I. Jordan, University of California, Berkeley
Michael Kearns, University of Pennsylvania
Scott Kirkpatrick, Hebrew University, Jerusalem
Daphne Koller, Stanford University
John Lafferty, University of Chicago
Todd K. Leen, Oregon Health & Sciences University
Richard Lippmann, Massachusetts Institute of Technology
Bartlett Mel, University of Southern California
John Moody, Intl. Computer Science Inst., Berkeley and Portland
John C. Platt, Microsoft Research
Gerald Tesauro, IBM Watson Labs
Sebastian Thrun, Stanford University
Dave Touretzky, Carnegie Mellon University
Lawrence Saul, University of California, San Diego
Bernhard Schölkopf, Max Planck Institute for Intelligent Systems,
Dale Schuurmans, University of Alberta, Canada
Sara A. Solla, Northwestern University Medical School

PROGRAM COMMITTEE

Jacob Abernethy (University of Michigan)
Ryan Adams (Harvard University)
Deepak Agarwal (LinkedIn)
Edo Airoldi (Harvard University)
Anima Anandkumar (UC Irvine)
Paolo Avesani (FBK - NiLab, Trento, Italy)
Borja Balle (McGill University)
Peter Bartlett (Queensland University/UC Berkeley)
Mikhail Belkin (Ohio State University)
Samy Bengio (Google)
Kristin Bennett (Rensselaer Polytechnic Institute)
Jeff Bilmes (University of Washington)
Matthew Blaschko (Inria Saclay)
David Blei (Princeton)
Karsten Borgwardt (MPI for Intelligent Systems)
Leon Bottou (Microsoft Research)
Guillaume Bouchard (Xerox Research Centre)
Michael Bowling (Alberta)
Sebastien Bubeck (Princeton)
Lawrence Carin (Duke)
Miguel Carreira-Perpinan (UC Merced)
Trevor Cohn (University of Melbourne)
Marco Cuturi (Kyoto University)
Florence d'Alche-Buc (Université d'Evry-Val d'Essonne)
Viriginia De Sa (UCSD)
Arnaud Doucet (Oxford)
François Fleuret (Idiap Research Institute)
Claudio Gentile (DiSTA Universita dell'Insubria)
Mohammad Ghavamzadeh (INRIA & Adobe Research)
Geoff Gordon (Carnegie Mellon)
Dilan Gorur (Yahoo! Labs)
Philipp Hennig (MPI Tuebingen)
Tom Heskes (Radboud)
Bert Huang (University of Maryland)
Christian Igel (University of Copenhagen)
Thorsten Joachims (Cornell)
Brian Kingsbury (IBM Research)
Jyrki Kivinen (University of Helsinki)
Andreas Krause (ETH Zurich)
Brian Kulis (Ohio State University)
Sanjiv Kumar (Google Research)
Gert Lanckriet (UCSD)
Hugo Larochelle (Université de Sherbrooke)
Guy Lebanon (Amazon)
Honglak Lee (University of Michigan)
Robert Legenstein (Graz University of Technology)
Lihong Li (Microsoft Research)
Phil Long (Microsoft)
Aurelie Lozano (IBM Research)
Lester Mackey (Stanford University)
Michael Mahoney (UC Berkeley)
Benjamin Marlin (UMass Amherst)
Mehryar Mohri (Courant Institute and Google Research)
Iain Murray (University of Edinburgh)
Hartmut Neven (Google)
Jennifer Neville (Purdue University)
Sebastian Nowozin (MSR Cambridge)
Guillaume Obozinski (Ecole des Ponts - ParisTech)
Mike Osborne (University of Oxford)
Sinno Pan (National Univ. of Singapore)
Jan Peters (TU Darmstadt)
Joaquin Quinonero-Candela (Facebook)
Marc'Aurelio Ranzato (Facebook)
Steve Renals (University of Edinburgh)
Afshin Rostamizadeh (Google Research)
Dale Schuurmans (University of Alberta)
D Sculley (Google Research)
Ricardo Silva (UCL)
Alex Slivkins (Microsoft Research)
Suvrit Sra (MPI for Intelligent Systems & CMU)
Oliver Stegle (EMBL-EBI)
Amos Storkey (University of Edinburgh)
Charles Sutton (University of Edinburgh)
Umar Syed (Google Research)
Bill Triggs (CNRS)
Koji Tsuda (Advanced Industrial Science and Technology)
Nicolas Usunier (Université de Technologie de Compiègne)
Eleni Vasilaki (University of Sheffield)
Sergei Vassilvitskii (Google)
Aki Vehtari (Aalto University)
Jean-Philippe Vert (Mines ParisTech)
Kilian Weinberger (Washington University)
Yair Weiss (Hebrew University)
Rich Wilkinson (University of Nottingham)
Robert Williamson (NICTA)
David Wipf (Microsoft Research)
Jennifer Wortman Vaughan (Microsoft Research)
John Wright (Columbia University)
Byron Yu (CMU)
Alan Yuille (UCLA)
Richard Zemel (University of Toronto)
Tong Zhang (Baidu/Rutgers)

SPONSORS

NIPS gratefully acknowledges the generosity of those individuals and organizations who have provided financial support for the NIPS 2014 conference. Their financial support enables us to sponsor student travel and participation, the outstanding paper awards, and the volunteers who assist during NIPS.

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.

Amazon.com (jobs: http://www.amazon.jobs/) 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 fut re 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.

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.

Ketchum Trading, LLC is a privately held, proprietary trading firm in Chicago. Driven by algorithms and utilizing the most advanced software and technology in the industry, our multi-strategy firm makes markets and trades in futures, options, cash equities, and exchange traded funds. The fundamentals of Ketchum Trading are simple: manage the risks inherent in modern markets, leverage state-of-the-art technology and proprietary models, and gain a mathematical advantage. The rest takes care of itself.

Artificial Intelligence Journal (AIJ) which commenced publication in 1970, is now the generally accepted premier international forum for the publication of results of current research in this field. The journal welcomes foundational and applied papers describing mature work involving computational accounts of aspects of intelligence. Specifically, it welcomes papers on: AI and Philosophy, automated reasoning and inference, case-based reasoning, cognitive aspects of AI, commonsense reasoning, constraint processing, heuristic search, high-level computer vision, intelligent interfaces, intelligent robotics, knowledge representation, machine learning, multi-agent systems, natural language processing, planning and theories of action, reasoning under uncertainty or imprecision. The journal reports results achieved; proposals for new ways of looking at AI problems must include demonstrations of effectiveness. Papers describing systems or architectures integrating multiple technologies are welcomed. AIJ also invites papers on applications, which should describe a principled solution, emphasize its novelty, and present an in-depth evaluation of the AI techniques being exploited. The journal publishes an annual issue devoted to survey articles and also hosts a “competition section” devoted to reporting results from AI competitions. From time to time, there are special issues devoted to a particular topic; such special issues always have open calls.

QUALCOMM

Qualcomm Incorporated is the world leader in 3G, 4G and next-generation wireless technologies. Qualcomm innovations are enabling ultra-personal mobile devices; shaping relevant, next-generation mobile experiences; and inspiring transformative new business models and services. Dedicated to accelerating mobile growth and progress worldwide, Qualcomm is transforming the way people live, learn, work and play. Headquartered in San Diego, California, Qualcomm is included in the S&P 500 Index and is a FORTUNE 500(r) company traded on the NASDAQ Stock Market(r) under the ticker symbol QCOM. Qualcomm Research and Development, a division of Qualcomm Technologies, Inc., is where many of the industry's most talented engineers and scientists come to collaborate, share innovative ideas, and create the wireless technologies and solutions that will transform the future of wireless.

Baidu Research

Baidu Research, based in Silicon Valley and Beijing, is led by Dr. Andrew Ng, Chief Scientist. Baidu Research comprises three interrelated labs: the Silicon Valley AI Lab, the Institute of Deep Learning and the Big Data Lab, led by Dr. Adam Coates, Dr. Kai Yu and Dr. Tong Zhang, respectively. The organization brings together global research talent to work on fundamental technologies in areas such as image recognition and image-based search, voice recognition, natural language processing and semantic intelligence.

United Technologies Research Center

United Technologies Research Center (UTRC) 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 open roles for people with strong machine learning, cloud analytics and HCI skills. If you're technically strong and enjoy continual learning of new application domains, UTRC is the place for you.

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.

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.

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.

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

NVIDIA

NVIDIA is a world leader in visual computing. Our technologies are transforming a world of displays into a world of interactive discovery for everyone from gamers and scientists, to consumers and enterprises. We invest in our people, our technologies, and our research and development efforts to deliver the highest quality products to customers across the globe. NVIDIA's culture inspires our team of world-class engineers and developers to be at the top of their game. Data scientists in both industry and academia use GPUs for machine learning to make groundbreaking improvements across a variety of applications including image classification, video analytics, speech recognition and natural language processing. With thousands of computational cores GPUs have become the processor of choice for processing big data for data scientists.

SPONSORS

As part of The Walt Disney Company, Disney Research draws on a legacy of innovation and technology leadership that continues to this day. In 1923, Walt Disney sold his animated/live-action series the Alice Comedies, founded his eponymous company, and launched a succession of firsts: The first cartoon with fully synchronized sound (1928). The first full-color cartoon (1932). The first animated feature film (1937). The first modern theme park (1955). The Walt Disney Company was also an early leader in entertainment technology development with inventions like the multiplane camera, Audio-Animatronics, Circle-Vision 360°, and Fantasound. In 2006, The Walt Disney Company acquired Pixar Animation Studios, a move that brought a host of valuable creative and technology assets, including a strong culture of excellence in research. Pixar is a major generator and publisher of world-class research in computer graphics. Its scientists contribute directly to Pixar's critically acclaimed films, consistently winning multiple technical Academy Awards®. The Pixar acquisition was a source of inspiration for the formation of Disney Research, and continues to influence the way we're organized and run. Disney Research was launched in 2008 as an informal network of research labs that collaborate closely with academic institutions such as Carnegie Mellon University and the Swiss Federal Institute of Technology Zürich (ETH). We're able to combine the best of academia and industry: we work on a broad range of commercially important challenges, we view publication as a principal mechanism for quality control, we encourage engagement with the global research community, and our research has applications that are experienced by millions of people. We're honoring Walt Disney's legacy of innovation by researching novel technologies and deploying them on a global scale.

Booking.com was established in 1996 in Amsterdam, The Netherlands and has grown into the world's leading online hotel and accommodation reservations company . We are part of the Priceline Group (NASDAQ: PCLN). Booking.com is still operated from Amsterdam and supported by over 130 offices in over 50 countries. The team 7.000+ people right across the world are dedicated to serving Booking.com's customers and property-owning partners. Our IT teams take care of creating a perfect shopping experience for our customers. Over 100 nationalities are employed at Booking.com. We are a melting pot of people from all over the world working together. Employees are encouraged to develop their own talents. We have a variety of trainings and other methods within our Booking.com University program to help to do just that. Our mission is to help leisure and business travellers, whatever their budgets, easily discover, book, and enjoy the world's best places to stay.Our values describe the "How" of Booking.com's operations: how we behave, how we deal with colleagues, customers, and partners. It basically describes who we are and what drives our behavior.

Winton Capital Management is a world leading quantitative investment manager specialising in applying advanced scientific techniques to the analysis of financial markets. The company was founded in 1997 and now employs over 300 professionals across offices in London, Oxford, Zurich, Hong Kong, New York, Tokyo and Sydney.

The Mission of Oracle Labs is straightforward: Identify, explore, and transfer new technologies that have the potential to substantially improve Oracle's business. Oracle's commitment to R&D is a driving factor in the development of technologies that have kept Oracle at the forefront of the computer industry. Although many of Oracle's leading-edge technologies originate in its product development organizations, Oracle Labs is the sole organization at Oracle that is devoted exclusively to research. The acquisition of Sun Microsystems, along with dozens of other acquired companies, brought a wide array of technologies to Oracle's portfolio. Oracle executives recognized that in Sun Microsystems Laboratories, Sun brought the combined company the benefits of an independent research organization - now renamed Oracle Labs.

AdRoll's Data Science team is responsible for continuously and automatically leveraging our platform's data to provide actionable insights and models for our business. Our critical mission is to decide the value of displaying our clients' ads to maximize their return on investment. Our proprietary large-scale machine learning infrastructure ingests billions of events with billions of sparse features every few hours to produce models queried thousands of times per second. This scale pushes us to engineer supervised and unsupervised solutions that are frequently beyond the scope of the literature

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

Superfish is a developer and provider of an innovative, next generation visual discovery platform that's out to transform the way images are searched, seen, utilized and shared over the web. Through its cutting-edge, patented technology and pioneering vision, Superfish has developed a cunningly intuitive visual search engine that analyzes images algorithmically, without the need of text tags or human intervention. The company is reinventing image search and focusing on empowering consumers and businesses to experience new ways to engage with and search visual content. As the social web has exploded and photos play such a large role in people's lives online, Superfish sees a huge opportunity in solving new visual search problems. Founded in 2006 by Adi Pinhas and Michael Chertok, Superfish is headquartered in Palo Alto, California and is backed by Draper Fisher Jurvetson (DFJ) and Vintage Investment Partners

BOSCH

Intellectual Inspiration. Entrepreneurial Freedom. Global Opportunities. It's all yours with a career at Bosch. "Made by Bosch" stands for the first-class quality of a global player. The Bosch Research and Technology Center (RTC), with labs in Palo Alto, CA, Pittsburgh, PA, and Cambridge, MA focus on innovative research and development. The Bosch Research and Technology Center focuses on topics such as algorithms for Robotics, Autonomous Systems and Data Mining. Our team of research scientists and software engineers in Palo Alto will grow rapidly in the next couple of years. Therefore, now is the perfect time for you to join and make an impact with your passion to innovate! Please check our website for our current openings: www.boschresearch.com

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.

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 the Xerox services offerings. Today 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.

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.

Symantec™

Symantec Corporation (NASDAQ: SYMC) is an information protection expert that helps people, businesses and governments seeking the freedom to unlock the opportunities technology brings – anytime, anywhere. Founded in April 1982, Symantec, a Fortune 500 company, operating one of the largest global data-intelligence networks, has provided leading security, backup and availability solutions for where vital information is stored, accessed and shared. The company's more than 21,500 employees reside in more than 50 countries. Ninety-nine percent of Fortune 500 companies are Symantec customers. In fiscal 2013, it recorded revenues of $6.9 billion. To learn more go to www.symantec.com or connect with Symantec at: go.symantec.com/socialmedia

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 support the NIPS Best Student Awards.

Frontiers in
Distributed Information
Systems Foundation

The focus is on the rapidly changing possibilities in the backbone and transport layers of the world's digital networks, and on the new services, new clients, and new communities that are coming into being as a result of their continuing growth. Opportunities and challenges are presented by high-performance computing and the explosion of digital storage.

PROGRAM HIGHLIGHTS

FRIDAY, DECEMBER 12TH

Breakfast - 7:00 - 8:30 Level 7

Registration Level 5
7:30 AM - Noon & 2:30 - 6:00 PM

FRIDAY WORKSHOPS

All workshops run from 8:30 - Noon & 3:00 - 6:30 PM
Coffee breaks 10:00 - 10:30 AM and 4:30 - 5:00 PM

ABC In Montreal
Level 5; room 512 a,e

Analyzing the Omics of the Brain
Level 5; room 515 a

Autonomously Learning Robots
Level 5; room 512 d, h

Bayesian Optimization in Academia and Industry
Level 5; room 512 b, f

Challenges in Machine Learning Workshop (CiML 2014)
Level 5; room 511c

Deep Learning and Representation Learning
Level 5; room 511a,b,d,e

Distributed Machine Learning and Matrix Computations
Level 5; room 510a

Fairness, Accountability, and Transparency in Machine Learning
Level 5; room 514c

From Bad Models to Good Policies (Sequential Decision Making under Uncertainty)
Level 5; room 512c,g

Large Scale Optical Physiology: From Data-acquisition to Models of Neural Coding
Level 5; room 514a,b

Learning Semantics
Level 5; room 513a,b

Machine Learning for Clinical Data Analysis, Healthcare and Genomics
Level 5; room 511f

Modern Machine Learning and Natural Language Processing
Level 5; room 510c

NIPS Workshop on Transactional Machine Learning and E-Commerce
Level 5; room 510d

OPT2014: Optimization for Machine Learning
Level 5; room 513e,f

Out of the Box: Robustness in High Dimension
Level 5; room 510b

Personalization: Methods and Applications
Level 5; room 513c,d

Perturbations, Optimization, and Statistics
Level 5; room 515b,c

SATURDAY, DECEMBER 13TH

Breakfast - 7:00 - 8:30 Level 7

Registration Level 5
7:30 AM - Noon

SATURDAY WOKSHOPS

All workshops run from 8:30 - Noon & 3:00 - 6:30 PM
Coffee breaks 10:00 - 10:30 AM and 4:30 - 5:00 PM

3rd NIPS Workshop on Probabilistic Programming
Level 5; room 512d,h

4th Workshop on Automated Knowledge Base Construction (AKBC)
Level 5; room 515

Advances in Variational Inference
Level 5; room 510a

Analysis of Rank Data: Confluence of Social Choice, Operations Research, and Machine Learning
Level 5; room 513c,d

Discrete Optimization in Machine Learning
Level 5; room 514

High-energy Particle Physics, Machine Learning, and the HiggsML Data Challenge (HEPML)
Level 5; room 511c

Human Propelled Machine Learning
Level 5; room 511f

Large-scale Reinforcement Learning and Markov Decision Problems
Level 5; room 511d

Machine Learning in Computational Biology
Level 5; room 510b

MLINI 2014 - 4th NIPS Workshop on Machine Learning and Interpretation in Neuroimaging: Beyond the Scanner
Level 5; room 513e,f

Modern Nonparametrics 3: Automating the Learning Pipeline
Level 5; room 511e

Networks: From Graphs to Rich Data
Level 5; room 510c

NIPS'14 Workshop on Crowdsourcing and Machine Learning
Level 5; room 511a

Novel Trends and Applications in Reinforcement Learning
Level 5; room 512a,e

Optimal Transport and Machine Learning
Level 5; room 512c,g

Representation and Learning Methods for Complex Outputs
Level 5; room 512b,f

Riemannian Geometry in Machine Learning, Statistics and Computer Vision
Level 5; room 511b

Second Workshop on Transfer and Multi-Task Learning: Theory meets Practice
Level 5; room 510d

Software Engineering for Machine Learning
Level 5; room 513a,b

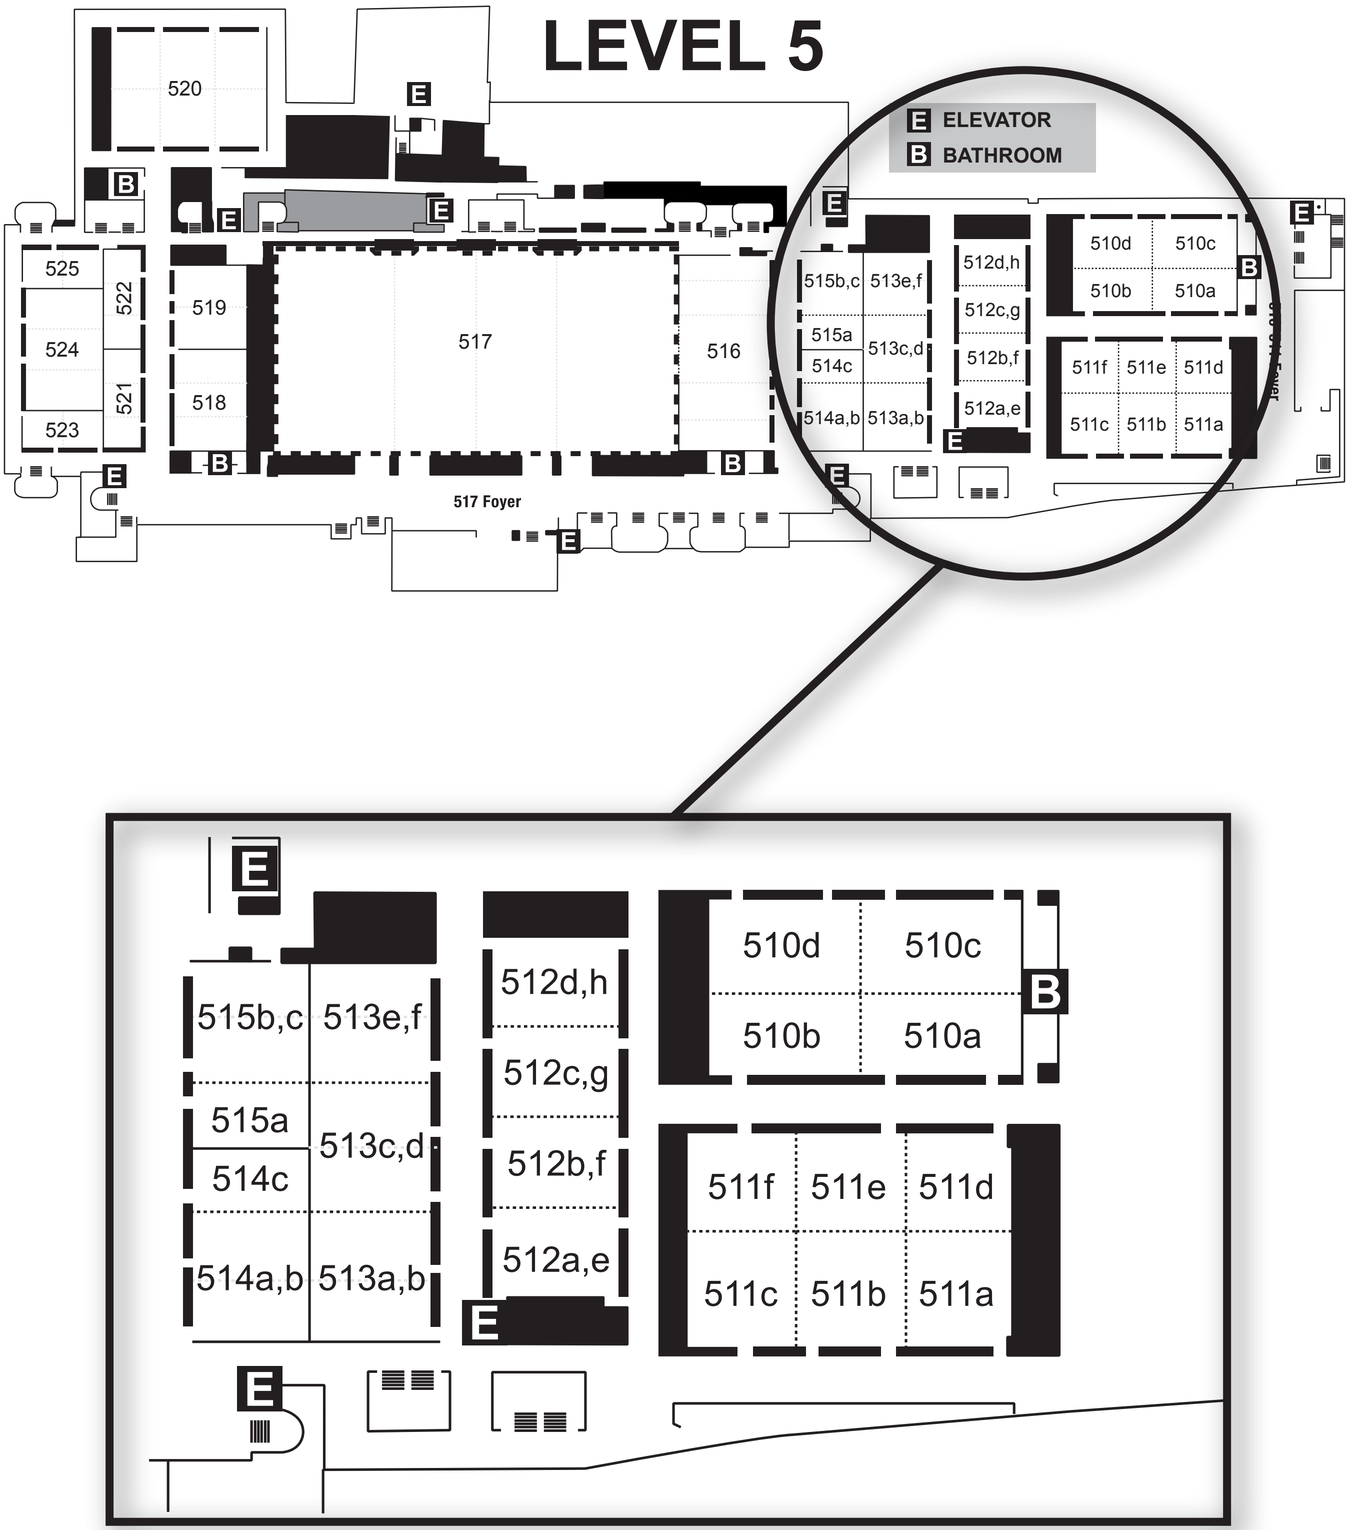
LEVEL 5
E ELEVATOR
B BATHROOM
520
525
524
523
522
521
519
518
517
516
517 Foyer
515b,c
513e,f
515a
514c
513c,d
514a,b
513a,b
512d,h
512c,g
512b,f
512a,e
510d
510c
510b
510a
511f
511e
511d
511c
511b
511a

FRIDAY WORKSHOPS

FRIDAY WORKSHOP
8:30AM - 6:30PM

ABC in Montreal

ORGANIZERS:

Ted Meeds	University of Amsterdam
Max Welling	University of Amsterdam
Richard Wilkinson	The University of Nottingham
Xi'an Robert	Université Paris-Dauphine
Neil Lawrence	University of Sheffield

LOCATION: Level 5; room 512a,e

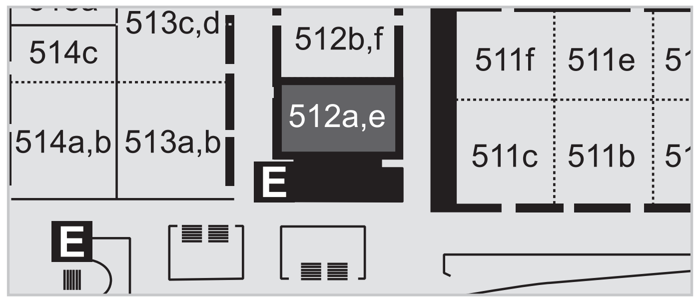

ABSTRACT:

Approximate Bayesian computation (ABC) or likelihood-free (LF) methods have developed mostly beyond the radar of the machine learning community, but are important tools for a large segment of the scientific community. This is particularly true for systems and population biology, computational psychology, computational chemistry, computational finance, etc. Recent work has applied both machine learning models and algorithms to general ABC inference (e.g., NN, forests, GPs, LDA) and ABC inference to machine learning (e.g. using computer graphics to solve computer vision using ABC). In general, however, there is significant room for more intense collaboration between both communities. Submissions on the following topics are encouraged (but not limited to):

Examples of topics of interest in the workshop include (but are not limited to):

* Applications of ABC to machine learning, e.g., computer vision, other inverse problems (RL)…
* ABC Reinforcement Learning (other inverse problems)
* Machine learning models of simulations, e.g., NN models of simulation responses, GPs etc.
* Selection of sufficient statistics and massive dimension reduction methods
* Online and post-hoc error
* ABC with very expensive simulations and acceleration methods (surrogate modeling, choice of design/simulation points)
* Relation between ABC and probabilistic programming
* Posterior evaluation of scientific problems/interaction with scientists
* Post-computational error assessment
* Impact on resulting ABC inference

WEBSITE: https://sites.google.com/site/abcinmontreal

Analyzing the Omics of the Brain

ORGANIZERS:

Gal Chechik	Google, BIU
Mark Reimers	VCU
Michael Hawrylycz	Allen Institute for Brain Science

LOCATION: Level 5; room 515a

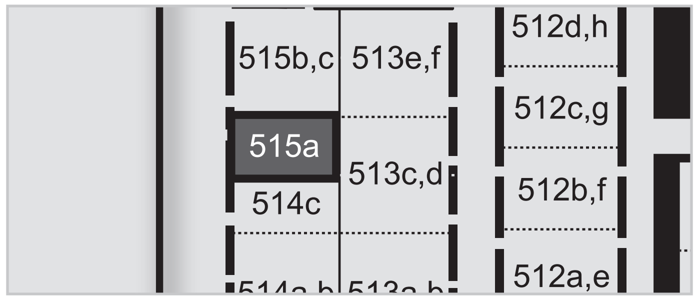

ABSTRACT:

In the past few years, the field of molecular biology of the brain has been transformed from hypothesis-based experiments to high-throughput experiments. The massive growth of data, including measures of the brain transcriptome, methylome and proteome, now raises new questions in neurobiology and new challenges in analysis of these complex and vast datasets. While many of these challenges are shared with other computational biology studies, the complexity of the brain poses special challenges. Brain genomics data includes high-resolution molecular imagery, developmental time courses and most importantly, underlies complex behavioral phenotypes and psychiatric diseases. New methods are needed to address questions about the brain-wide, genome-wide and life-long genomic patterns in the brain and their relation to brain functions like plasticity and information processing.

The goal of the workshop is to bring together people from the neuroscience, cognitive science and the machine learning community. It aims to ease the path for scientists to connect the wealth of genomic data to the issues of cognition and learning that are central to NIPS, with an eye to the emerging high-throughput behavioral data which many are gathering. We invite contributed talks on novel methods of analysis to brain genomics, as well as techniques to make meaningful statistical relationships to phenotypes.

The target audience includes two main groups: people interested in developing machine learning approaches to neuroscience, and people from neuroscience and cognitive science interested in connecting their work to brain genomics.

WEBSITE: http://chechiklab.biu.ac.il/~gal/BrainOmics2014/

Autonomously Learning Robots

ORGANIZERS:

Gerhard Neumann	TU Darmstadt
Joelle Pineau	McGill University
Peter Auer	University of Leoben
Marc Toussaint	Universty Suttgart

LOCATION: Level 5; room 512 d,h

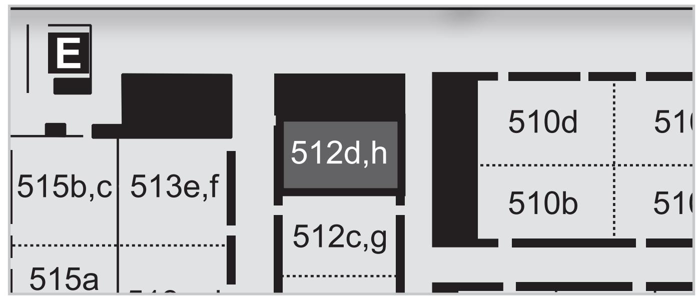

ABSTRACT:

To autonomously assist human beings, future robots have to autonomously learn a rich set of complex behaviors. So far, the role of machine learning in robotics has been limited to solve pre-specified sub-problems that occur in robotics and, in many cases, off-the-shelf machine learning methods. The approached problems are mostly homogeneous, e.g., learning a single type of movement is sufficient to solve the task, and do not reflect the complexities that are involved in solving real-world tasks.

In a real-world environment, learning is much more challenging than solving such homogeneous problems. The agent has to autonomously explore its environment and discover versatile behaviours that can be used to solve a multitude of different tasks throughout the future learning progress. It needs to determine when to reuse already known skills by adapting, sequencing or combining the learned behaviour and when to learn new behaviours. To do so, it needs to autonomously decompose complex real-world tasks into simpler sub-tasks such that the learned solutions for these sub-tasks can be re-used in a new situation. It needs to form internal representations of its environment, which is possibly containing a large variety of different objects or also different agents, such as other robots or humans. Such internal representations also need to shape the structure of the used policy and/or the used value function of the algorithm, which need to be flexible enough such to capture the huge variability of tasks that can be encountered in the real world. Due to the multitude of possible tasks, it also cannot rely on a manually tuned reward function for each task, and, hence, it needs to find a more general representations for the reward function. Yet, an autonomous robot is likely to interact with one or more human operators that are typically experts in a certain task, but not necessarily experts in robotics. Hence, an autonomously learning robot also should make effective use of feedback that can be acquired from a human operator. Typically, different types of instructions from the human are available, such as demonstrations and evaluative feedback in form of a continuous quality rating, a ranking between solutions or a set of preferences. In order to facilitate the learning problem, such additional human instructions should be used autonomously whenever available. Yet, the robot also needs to be able to reason about its competence to solve a task. If the robot thinks it has poor competence or the uncertainty of the competence is high, the robot should request more instructions from the human expert.

Most machine learning algorithms are missing these types of autonomy. They still rely on a large amount of engineering and fine-tuning from a human expert. The human typically needs to specify the representation of the reward-function, of the state, of the policy or of other internal representations used by the learning algorithms. Typically, the decomposition of complex tasks into sub-tasks is performed by the human expert and the parameters of such algorithms are fine tuned by hand. The algorithms typically learn from a pre-specified source of feedback and can not autonomously request more instructions such as demonstrations, evaluative feedback or corrective actions. We belief that this lack of autonomy is one of the key reasons why robot learning could not be scaled to more complex, real world tasks. Learning such tasks would require a huge amount of fine tuning which is very costly on real robot systems.

WEBSITE: http://www.ias.tu-darmstadt.de/Workshops/NIPS2014

FRIDAY WORKSHOP
8:30AM - 6:30PM

Bayesian Optimization in Academia and Industry

ORGANIZERS:

Matthew Hoffman	University of Cambridge
Zoubin Ghahramani	University of Cambridge
Ryan Adams	Harvard University
Jasper Snoek	Harvard University
Kevin Swersky	University of Toronto

LOCATION: Level 5; room 512 b,f

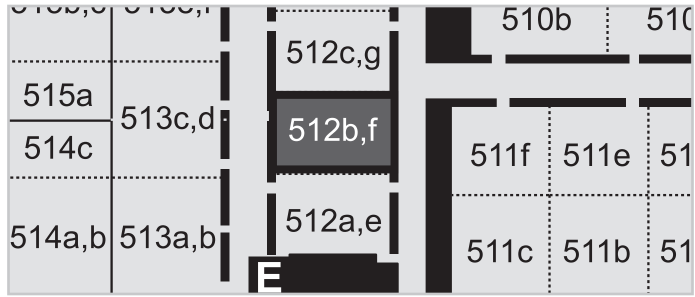

ABSTRACT:

Bayesian optimization has emerged as an exciting subfield of machine learning that is concerned with the global optimization of noisy, black-box functions using probabilistic methods. Systems implementing Bayesian optimization techniques have been successfully used to solve difficult problems in a diverse set of applications. There have been many recent advances in the methodologies and theory underpinning Bayesian optimization that have extended the framework to new applications as well as provided greater insights into the behaviour of these algorithms. Bayesian optimization is now increasingly being used in industrial settings, providing new and interesting challenges that require new algorithms and theoretical insights.

At last year's NIPS workshop on Bayesian optimization the focus was on the intersection of "Theory and Practice". The workshop this year will follow this trend by again looking at theoretical contributions, but also by focusing on the practical side of Bayesian optimization in industry. The goal of this workshop is not only to bring together both practical and theoretical research knowledge from academia, but also to facilitate cross-fertilization with industry. Specifically, we would like to carefully examine the types of problems where Bayesian optimization works well in industrial settings, but also the types of situations where additional performance is needed. The key questions we will discuss are: how to scale Bayesian optimization to long time-horizons and many observations? How to tackle high-dimensional data? How to make Bayesian optimization work in massive, distributed systems? What kind of structural assumptions are we able to make? And finally, what can we say about these questions both empirically and theoretically?

The target audience for this workshop consists of both industrial and academic practitioners of Bayesian optimization as well as researchers working on theoretical advances in probabilistic global optimization. To this end we have invited many industrial users of Bayesian optimization to attend and speak at the workshop. We expect this exchange of industrial and academic knowledge will lead to a significant interchange of ideas and a clearer understanding of the challenges and successes of Bayesian optimization as a whole.

A further goal is to encourage collaboration between the diverse set of researchers involved in Bayesian optimization. This includes not only interchange between industrial and academic researchers, but also between the many different sub-fields of machine learning which make use of Bayesian optimization. We are also reaching out to the wider global optimization and Bayesian inference communities for involvement.

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

Challenges in Machine Learning Workshop (CiML 2014)

ORGANIZERS:

Isabelle Guyon	Clopinet
Evelyne Viegas	Microsoft Research

LOCATION: Level 5; room 511c

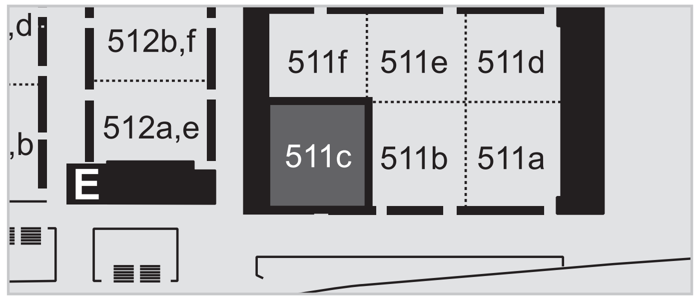

ABSTRACT:

Challenges in Machine Learning have proven to be efficient and cost-effective ways to quickly bring to industry solutions that may have been confined to research. In addition, the playful nature of challenges naturally attracts students, making challenge a great teaching resource. Challenge participants range from undergraduate students to retirees, joining forces in a rewarding environment allowing them to learn, perform research, and demonstrate excellence. Therefore challenges can be used as a means of directing research, advancing the state-of-the-art or venturing in completely new domains.

Yet, despite initial successes and efforts made to facilitate challenge organization with the availability of competition platforms, little effort has been put into the theoretical foundations of challenge design and the optimization of challenge protocols. This workshop will bring together workshop organizers, platform providers, and participants to discuss best practices in challenge organization and new methods and application opportunities to design high impact challenges. The themes to be discussed will include new paradigms of challenge organization to tackle complex problems (e.g. tasks involving multiple data modalities and/or multiple levels of processing).

WEBSITE: http://ciml.chalearn.org

FRIDAY WORKSHOP
8:30AM - 6:30PM

Deep Learning and Representation Learning

ORGANIZERS:

Yoshua Bengio	University of Montreal
Roland Memisevic	University of Montreal
Adam Coates	Baidu Research
Andrew Ng	Stanford University

LOCATION: Level 5; room 511a,b,d,e

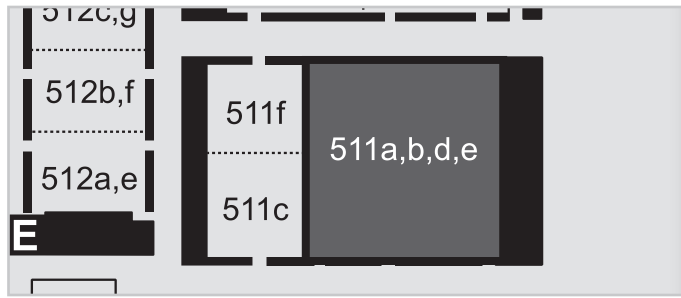

ABSTRACT:

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: http://www.dlworkshop.org/

FRIDAY WORKSHOP
8:30AM - 6:30PM

Distributed Machine Learning and Matrix Computations

ORGANIZERS:

Reza Bosagh Zadeh	Stanford University
Ion Stoica	UC Berkeley
Ameet Talwalkar	UC Berkeley

LOCATION: Level 5; room 510a

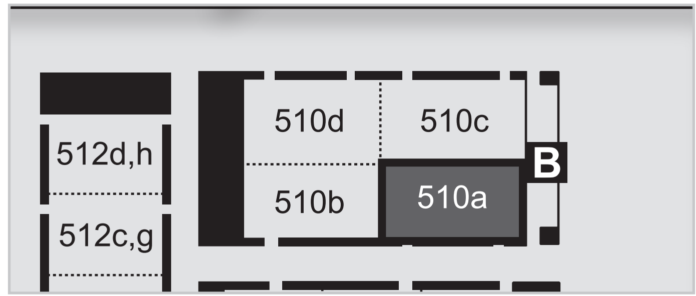

ABSTRACT:

The emergence of large distributed matrices in many applications has brought with it a slew of new algorithms and tools. Over the past few years, machine learning and numerical linear algebra on distributed matrices has become a thriving field. Manipulating such large matrices makes it necessary to think about distributed systems issues such as communication cost. This workshop aims to bring closer researchers in distributed systems and large scale numerical linear algebra to foster cross-talk between the two fields. The goal is to encourage distributed systems researchers to work on machine learning and 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://snap.stanford.edu/networks2013/

Fairness, Accountability, and Transparency in Machine Learning

ORGANIZERS:

Moritz Hardt — IBM Research Almaden
Solon Barocas — Princeton University

LOCATION: Level 5; room 514c

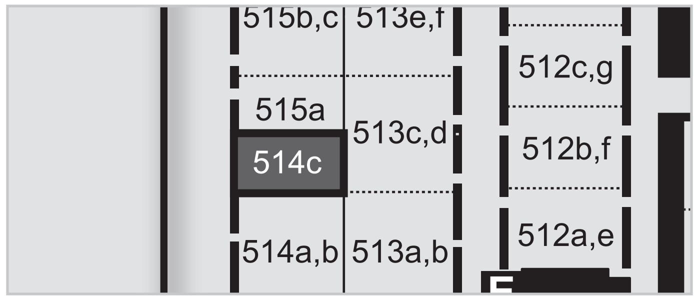

ABSTRACT:

We propose an interdisciplinary workshop on fairness in machine learning. As its role in decision-making in such areas as policing, employment, health, education, and commerce has grown, machine learning has raised corresponding concerns about the fairness of these decisions.

Reflecting these concerns, President Obama at the start of 2014 instructed the White House Office of Science and Technology Policy and the President's Council of Advisors for Science and Technology to undertake a 90-day review of Big Data. The resulting report, "Big Data: Seizing Opportunities, Preserving Values" [1], concluded that "big data technologies can cause societal harms beyond damages to privacy," and stressed that algorithms employed for eligibility determinations must be scrutinized for potentially discriminatory effects, even absent discriminatory intent. In its recommendations to the President, the report called for additional "technical expertise to stop discrimination" and for further research into the dangers of "encoding discrimination in automated decisions".

Our workshop takes up this call. As core tasks in machine learning become matters of growing policy concern, we must bring together members of the computer science and machine learning community to address issues of fairness, accountability, and transparency. The workshop will focus on these issues both as challenging constraints on the practical application of machine learning, as well as problems that can lend themselves to novel computational solutions.

Questions to the machine learning community include:

- How can we achieve high classification accuracy while eliminating discriminatory biases? What are meaningful formal fairness properties?
- How can we design expressive yet easily interpretable classifiers?
- Can we ensure that a classifier remains accurate even if the statistical signal it relies on is exposed to public scrutiny?
- Are there practical methods to test existing classifiers for compliance with a policy?

Participants will work together to understand the key normative and legal issues at stake, map the relevant computer science scholarship, evaluate the state of the solutions thus far proposed, and explore opportunities for new research and thinking within machine learning itself.

Our confirmed speakers and panelists include research leaders from several relevant disciplines including machine learning, data mining, computational social science, information security, policy, and the law.

[1] http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf

WEBSITE: http://fatml.org/

FRIDAY WORKSHOP
8:30AM - 6:30PM

From Bad Models to Good Policies (Sequential Decision Making under Uncertainty)

ORGANIZERS:

Odalric-Ambrym Maillard	Technion
Timothy Mann	Technion
Shie Mannor	Technion
Jeremie Mary	INRIA / Univ. Lille
Laurent Orseau	AgroParisTech/INRA

LOCATION: Level 5; room 512 c,g

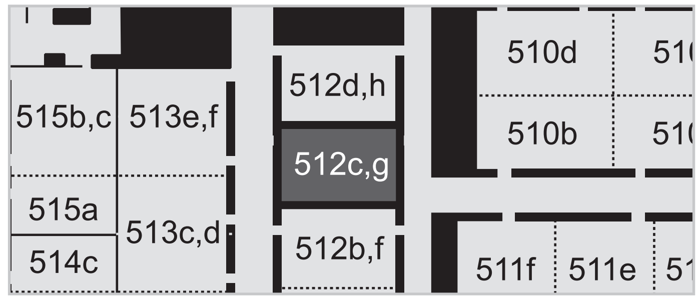

ABSTRACT:

This workshop aims to gather researchers in the area of sequential decision making to discuss recent findings and new challenges around the concept of model misspecification. A misspecified model is a model that either (1) cannot be tractably solved, (2) solving the model does not produce an acceptable solution for the target problem, or (3) the model clearly does not describe the available data perfectly. However, even though the model has its issues, we are interested in finding a good policy. The question is thus: How can misspecified models be made to lead to good policies?

We refer to the following (non exhaustive) types of misspecification.

1. States and Context. A misspecified state representation relates to research problems such as Hidden Markov Models, Predictive State Representations, Feature Reinforcement Learning, Partially Observable Markov Decision Problems, etc. The related question of misspecified context in contextual bandits is also relevant.

2. Dynamics. Consider learning a policy for a class of several MDPs rather than a single MDP, or optimizing a risk averse (as opposed to expected) objective. These approaches could be used to derive a reasonable policy for the target MDP even if the model we solved to obtain it is misspecified. Thus, robustness, safety, and risk-aversion are examples of relevant approaches to this question.

3. Actions. The underlying insight of working with high-level actions built on top of lower-level actions is that if we had the right high-level actions, we would have faster learning/planning. However, finding an appropriate set of high-level actions can be difficult. One form of model misspecification occurs when the given high-level actions cannot be combined to derive an acceptable policy.

More generally, since misspecification may slow learning or prevent an algorithm from finding any acceptable solution, improving the efficiency of planning and learning methods under misspecification is of primary importance. At another level, all these challenges can benefit greatly from the identification of finer properties of MDPs (local recoverability, etc.) and better notions of complexity. These questions are deeply rooted in theory and in recent applications in fields diverse as air-traffic control, marketing, and robotics. We thus also want to encourage presentations of challenges that provide a red-line and agenda for future research, or a survey of the current achievements and difficulties. This includes concrete problems like Energy management, Smart grids, Computational sustainability and Recommender systems.

We welcome contributions on these exciting questions, with the goals of (1) helping close the gap between strong theoretical guarantees and challenging application requirements, (2) identifying promising directions of near future research, for both applications and theory of sequential decision making, and (3) triggering collaborations amongst researchers on learning good policies despite being given misspecified models.

WEBSITE: https://sites.google.com/site/badmodelssdmuworkshop2014/

FRIDAY WORKSHOP
8:30AM - 6:30PM

Large Scale Optical Physiology: From Data-acquisition to Models of Neural Coding

ORGANIZERS:

Ferran Diego Andilla	University of Heidelberg
Jeremy Freeman	HHMI
Janelia Farm	MPI for Biological Cybernetics
Jakob Macke	MPI for Biological Cybernetics
Il Memming Park	UT Austin
Eftychios Pnevmatikakis	Columbia University

LOCATION: Level 5; room 514a,b

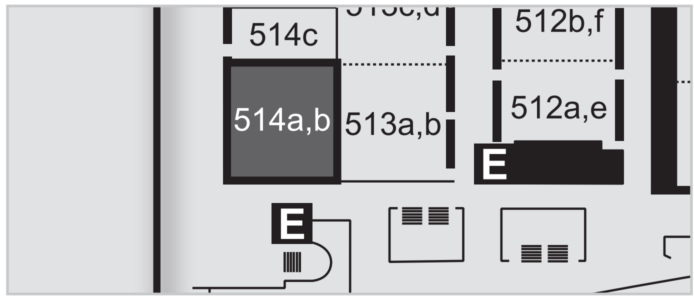

ABSTRACT:

A detailed understanding of brain function is a still-elusive grand challenge. Major advances in recording technologies (e.g. 2-photon and light-sheet microscopic imaging of calcium signals) are now beginning to provide measurements of neural activity at an unprecedented size and quality. Computational tools will be of critical importance both for the high-throughput acquisition and analysis of large-scale datasets. Reliable and robust tools for automated high-throughput analysis of such data that works have not been available so far. As a consequence, experimental reality is still characterized by semi-manual analysis or makeshift scripts that are specialized to a single setting. Similarly, many analysis still focus on the response properties of single neurons or on pairwise correlations across neurons, thereby potentially missing information which is only available at the population level.

The goal of this workshop is to discuss challenges and opportunities for computational neuroscience and machine learning which arise from large-scale recording techniques:

* 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? What are the computational bottlenecks for their analysis?

* What kind of computational tools play an important role on high-throughput data acquisition, e. g. visualization/ dimensionality reduction/information quantification? How can we figure out which algorithms work best, and which are the important challenges that are not met by existing techniques?

* What have we really learned from high-dimensional recordings that is new? What theories could we test, if only we had access to recordings from more neurons at the same time? What kind of statistics will be powerful enough to verify/falsify population coding theories? What can we infer about the network structure and dynamics?

We have invited scientists whose research addresses these questions, including researchers developing recording technologies, experimental and computational neuroscientists. We foresee active discussions amongst this multidisciplinary group of scientists to create a chance to discuss priorities and perspective, debate about the currently most relevant problems in the field, and emphasize the most promising future research directions. The target audience of this workshop includes industry and academic researchers interested in machine learning, neuroscience, big data and statistical inference.

WEBSITE: http://hci.iwr.uni-heidelberg.de//Staff/fdiego/LargeScaleOpticalPhysiology/

Learning Semantics

ORGANIZERS:

Cedric Archambeau	Amazon, Berlin
Antoine Bordes	Facebook
Leon Bottou	Microsoft Research
Chris Burges	Microsoft Research
David Grangier	Facebook Research

LOCATION: Level 5; room 513 a,b

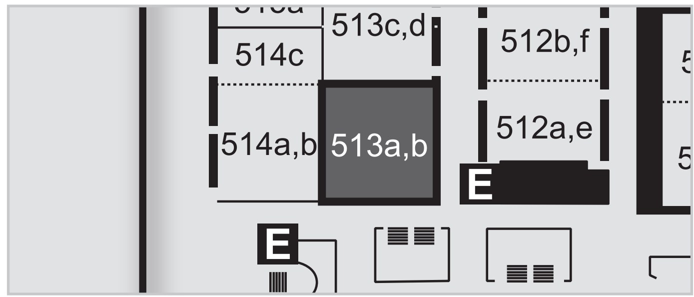

ABSTRACT:

Understanding the semantic structure of unstructured data -- text, dialogs, images -- is a critical challenge given their central role in many applications, including question answering, dialog systems, information retrieval... In recent years, there has been much interest in designing models and algorithms to automatically extract and manipulate these semantic representations from raw data.

Semantics is a diverse field. It encompases extracting structured data from text and dialog data (knowledge base extraction, logical form extraction, information extraction), linguistic approaches to extract and compose representation of meaning, inference and reasoning over meaning representation based on logic or algebra. It also includes approaches that aims at grounding language by learning relations between language and visual observations, linking language to the physical world (e.g. through robotics, machine commands). Despite spanning different disciplines with seemingly incompatible views, these approaches to semantics all aims at enabling computers to evolve and interact with humans and the physical world in general.

The goal of the workshop is dual. First, we aim at gathering experts from the different fields of semantics to favor cross-fertilization, discussions and constructive debates. Second, we encourage invited speakers and participants to expose their future research directions, take position and highlight the key challenges the community need to face. The workshop devotes most of the program to panel sessions about future directions.

**** Contributions ****

We will welcome contributions (up to 4 pages abstract) in the following areas and related topics:
- Word similarities and sense disambiguation
- Information and relation extraction
- Lexical and compositional semantics
- Learning semantic frames and semantic role labelling
- Grounded language learning
- Semantic representation for dialog understanding
- Visual scene understanding
- Multi-modal semantic representation and reasoning

WEBSITE: https://sites.google.com/site/learningsemantics2014/

FRIDAY WORKSHOP
8:30AM - 6:30PM

Machine Learning for Clinical Data Analysis, Healthcare and Genomics

ORGANIZERS:

Madalina Fiterau	Carnegie Mellon University
Julia Vogt	Sloan Kettering Institute
Gunnar Rätsch	Memorial Sloan-Kettering Cancer Center

LOCATION: Level 5; room 511f

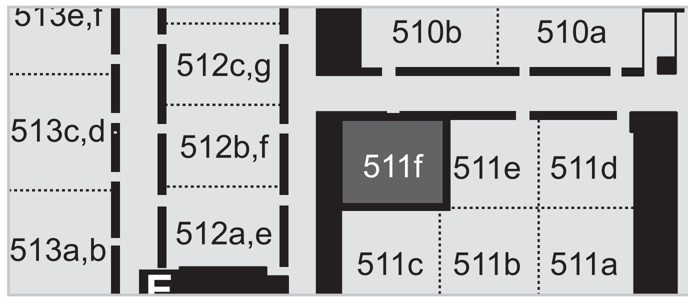

ABSTRACT:

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 labelling and clustering, clinical outcome prediction, efficient and scalable processing of medical records, feature selection or dimensionality reduction in clinical data, tools for personalized medicine, time-series analysis with medical applications and clinical genomics.

WEBSITE: http://goo.gl/gbfUl0

FRIDAY WORKSHOP
8:30AM - 6:30PM

Modern Machine Learning and Natural Language Processing

ORGANIZERS:

Ankur Parikh	Carnegie Mellon University
Avneesh Saluja	Carnegie Mellon University
Chris Dyer	Carnegie Mellon University
Eric Xing	Carnegie Mellon University

LOCATION: Level 5; room 510c

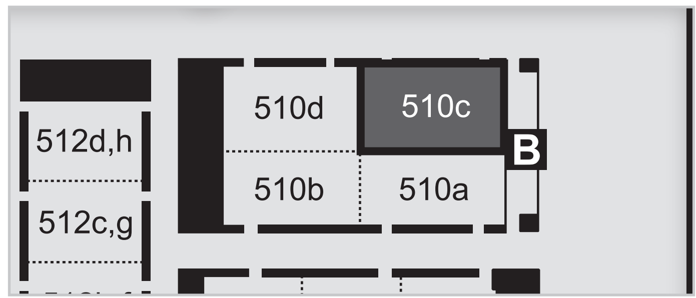

ABSTRACT:

The structure, complexity, and sheer diversity and variety of human language makes Natural Language Processing (NLP) distinct from other areas of AI. Certain core NLP problems have traditionally been an inspiration for machine learning (ML) solutions e.g., sequence tagging, syntactic parsing, and language modeling, primarily because these tasks can be easily abstracted into machine learning formulations (e.g., structured prediction, dimensionality reduction, or simpler regression and classification techniques). In turn, these formulations have facil- itated the transfer of ideas such as (but not limited to) discriminative methods, Bayesian nonparametrics, neural networks, and low rank / spectral techniques into NLP. Problems in NLP are particularly appealing to those doing core ML research due to the high-dimensional nature of the spaces involved (both the data and the label spaces) and the need to handle noise robustly, while principled, well-understood ML techniques are attractive to those in NLP since they potentially offer a solution to ill-behaved heuristics and training-test domain mismatch due to the lack of generalization ability these heuristics possess.

But there are many other areas within NLP where the ML community is less involved, such as semantics, discourse and pragmatics analysis, summarization, and parts of machine translation, and that continue to rely on linguistically- motivated but imprecise heuristics which may benefit from new machine learning approaches. Similarly, there are paradigms in ML, statistics, and optimization ranging from sub-modularity to bandit theory to Hilbert space embeddings that have not been well explored in the context of NLP.

The goal of this workshop is to bring together both applied and theoretical researchers in natural language processing and machine learning to facilitate the discussion of new frameworks that can help advance modern NLP. Some key questions we will address include (but not limited to):

- How can ML help provide novel representations and models to capture the structure of natural language?
- What NLP problems could benefit from new inference/ optimization techniques?
- How can we design new ML paradigms to address the lack of annotated data in complex structured prediction problems such as knowledge extraction and semantics?
- What technical challenges posed by multilinguality, lexical variation in social media, and nonstandard dialects are under-researched in ML?
- Does ML offer more principled ways for dealing with “overfitting” resulting from repeated evaluation on the same benchmark datasets?
- How can we tackle “scalability bottlenecks” unique to natural language?

Interest amongst both communities is high, as evidenced by a previous joint ACL-ICML symposium (2011) and a joint NAACL-ICML symposium (2013), and we hope to continue the exploration of topics beneficial to both fields that these symposiums initiated.

WEBSITE: http://www.cs.cmu.edu/~apparikh/nips2014ml-nlp/index.html

FRIDAY WORKSHOP
8:30AM - 6:30PM

NIPS Workshop on Transactional Machine Learning and E-Commerce

ORGANIZERS:

Amos Storkey — University of Edinburgh
Jacob Abernethy — University of Michigan
Mark Reid — The Australian National University

LOCATION: Level 5; room 510d

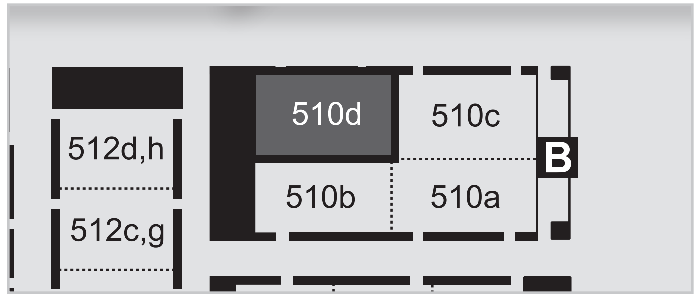

ABSTRACT:

In the context of building a machine learning framework that scales, the current modus operandi is a monolithic, centralised model building approach. These large scale models have different components, which have to be designed and specified in order to fit in with the model as a whole. The result is a machine learning process that needs a grand designer. It is analogous to a planned economy.

There is an alternative. Instead of a centralised planner being in charge of each and every component in the model, we can design incentive mechanisms for independent component designers to build components that contribute to the overall model design. Once those incentive mechanisms are in place, the overall planner need no longer have control over each individual component. This is analogous to a market economy.

The result is a transactional machine learning. The problem is transformed to one of setting up good incentive mechanisms that enable the large scale machine learning models to build themselves. Approaches of this form have been discussed in a number of different areas of research, including machine learning markets, collectives, agent-directed learning, ad-hoc sensor networks, crowdsourcing and distributed machine learning.

It turns out that many of the issues in incentivised transactional machine learning are also common to the issues that turn up in modern e-commerce setting. These issues include issues of mechanism design, encouraging idealised behaviour while modelling for real behaviour, issues surrounding prediction markets, questions of improving market efficiencies, and handling arbitrage, issue on matching both human and machine market interfaces and much more. On the theoretical side, there is a direct relationships between scoring rules, market scoring rules, and exponential family via Bregman Divergences. On the practical side, the issues that turn up in auction design relate to issues regarding efficient probabilistic inference.

The chances for each community to make big strides from understanding the developments in the others is significant. This workshop will bring together those involved in transactional and agent-based methods for machine learning, those involved in the development of methods and theory in e-commerce, those considering practical working algorithms for e-commerce or distributed machine learning and those working on financially incentivised crowdsourcing. The workshop will explore issues around incentivisation, handling combinatorial markets, and developing distributed machine learning. However the primary benefit will be the interaction and informal discussion that will occur throughout the workshop.

This topic is of particular interest because of the increasing importance of machine learning in the e-commerce setting, and the increasing interest in a distributed large scale machine learning. The workshop has some flavour of “multidisciplinary design optimization”: perhaps the optimum of the simultaneous problem of machine learning and e-commerce design is superior to the design found by optimizing each discipline sequentially, since it can exploit the interactions between the disciplines.

The expected outcomes are long lasting interactions between the communities and novel ideas in each individual community gained from learning from the others. The target group of participants are those working in machine learning markets, collectives, agent-directed learning, ad-hoc sensor networks, economic mechanisms in crowdsourcing and distributed machine learning, those working in areas of economics and markets, along with those looking at theory or practice in e-commerce, ad auctions, prediction markets and market design.

WEBSITE: http://workshops.inf.ed.ac.uk/ml/nipstransactional/

FRIDAY WORKSHOP
8:30AM - 6:30PM

OPT2014: Optimization for Machine Learning

ORGANIZERS:

Zaid Harchaoui	INRIA
Suvrit Sra	MPI for Intelligent Systems
Alekh Agarwal	Microsoft Research
Miro Dudik	Microsoft Research
Martin Jaggi	ETH Zurich
Aaditya Ramdas	Carnegie Mellon University

LOCATION: Level 5; room 513e,f

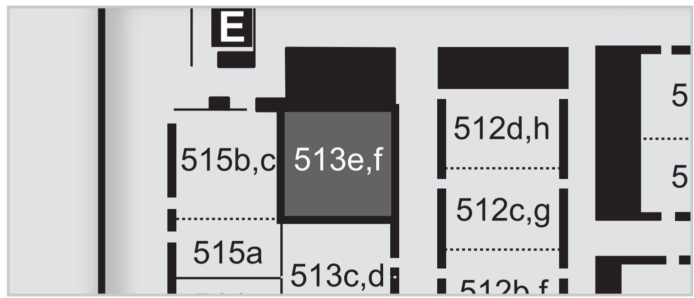

ABSTRACT:

As the seventh in its series, OPT 2014 stands on significant precedent established by OPT 2008--OPT 2013, 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 the 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.

This year we intend to add a new flavor to this successful series of NIPS workshops: special topics. In particular, we propose to structure both the invited talks and contributed talks around a current, core challenge in ML. The workshop will provide a great opportunity to gather renowned experts working towards similar goals, and will engender a fruitful exchange of ideas and discussions towards tackling this challenge.

For OPT2014, the special topic will be the “Curses and Blessings of Non-Convexity”, with attention to practical and theoretical advances in our understanding of difficult non-convex optimization. The three invited speakers will cover some of these advances (JB Lasserre, sum-of-squares, polynomial optimization, complexity), theoretical and methodological advances in non-convex optimization (A. Beck, results on alternating optimization), current challenges and new approaches for training deep architectures with non-convex optimization algorithms (Y Bengio, saddle-points in deep learning). The workshop will also welcome contributed talks and posters on traditional topics in optimization for machine learning.

For further details regarding the intended audience, list of topics covered, confirmed invited speakers and program committee, please refer to the web site at the bottom of this page

Organizational details not in the CFP are summarized here. The main features of the proposed workshop are:
1. One day long with morning and afternoon sessions
2. Three invited talks by leading experts from optimization and ML
3. Focus on discussion, which we plan to further promote by having an open problems session
4. Contributed talks from the broader OPT and ML community
5. An interactive poster session

Invited Speakers
Our workshop will feature experts from both optimization and ML. The tentative list of invited speakers for this year is: * Jean Bernard Lasserre (CNRS, France) * Yoshua Bengio (University of Montreal, Canada) * Amir Beck (Technion, Israel)

WEBSITE: http://www.opt-ml.org

FRIDAY WORKSHOP
8:30AM - 6:30PM

Out of the Box: Robustness in High Dimension

ORGANIZERS:

Stephen Becker — University of Colorado
Aleksandr Aravkin — IBM TJ Watson Research Center
Aurelie Lozano — IBM T.J. Watson Research Center

LOCATION: Level 5; room 510b

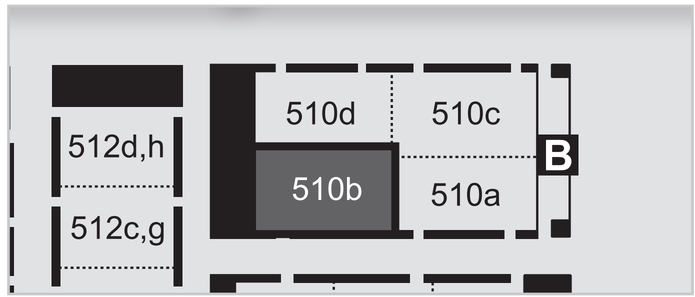

ABSTRACT:

The technical term “robust” was coined in 1953 by G. E. P. Box and exemplifies his adage, “all models are wrong, but some are useful”. Over the past decade, a broad range of new paradigms have appeared that allow useful inference when standard modeling assumptions are violated. Classic examples include heavy tailed formulations that mitigate the effect of outliers which would otherwise degrade the performance of Gaussian-based methods.

High-dimensional data are becoming ubiquitous in diverse domains such as genomics, neuroimaging, economics, and finance. Such data exacerbate the relevance of robustness as errors and model misspecification are prevalent in such modern applications. To extract pertinent information from large scale data, robust formulations require a comprehensive understanding of machine learning, optimization, and statistical signal processing, thereby integrating recovery guarantees, statistical and computational efficiency, algorithm design and scaling issues. For example, robust Principal Component Analysis (RPCA) can be approached using both convex and nonconvex formulations, giving rise to tradeoffs between computational efficiency and theoretical guarantees.

The goal of this workshop is to bring together machine learning, high-dimensional statistics, optimization and select large-scale applications, in order to investigate the interplay between robust modeling and computation in the large-scale setting.

We incorporate several important examples that are strongly linked by this theme:

(a) Low rank matrix recovery, robust PCA, and robust dictionary learning
(b) Robust inference for large scale inverse problems and machine learning
(c) Non-convex formulations: heavy tails, factorized matrix inversion, nonlinear forward models
(d) Robust optimization: avoiding overfitting on precise but unreliable parameters

It is the aim of this workshop to bring together researchers from statistics, machine learning, optimization, and applications, in order to focus on a comprehensive understanding of robust modeling and computation. In particular, we will see challenges of implementing robust formulations in the large-scale and nonconvex setting, as well as examples of success in these areas.

The workshop follows in the footsteps if the “Robust ML” workshop at NIPS in 2010. The field is very active and there have been significant advances in the past 4 years. We also expect to have new topics, such as new applications of robust optimization to user-perturbed problems and Markov Decision Processes.

WEBSITE: https://sites.google.com/site/saravkin/nips2014-robust

FRIDAY WORKSHOP
8:30AM - 6:30PM

Personalization: Methods and Applications

ORGANIZERS:

Khalid El-Arini	Facebook
Yisong Yue	California Institute of Technology
Dilan Gorur	Microsoft

LOCATION: Level 5; room 513c,d

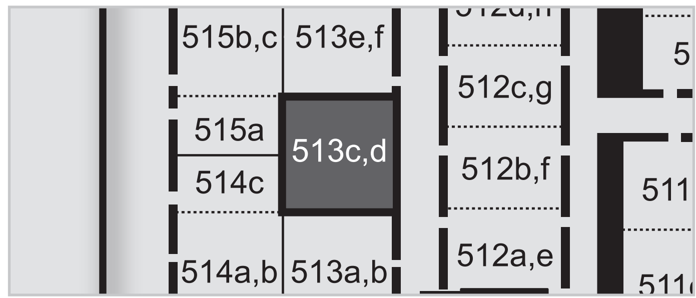

ABSTRACT:

From online news to online shopping to scholarly research, we are inundated with a torrent of information on a daily basis. With our limited time, money and attention, we often struggle to extract actionable knowledge from this deluge of data. A common approach for addressing this challenge is personalization, where results are automatically filtered to match the tastes and preferences of individual users.

This workshop aims to bring together researchers from industry and academia in order to describe recent advances and discuss future research directions pertaining to the personalization of digital systems, broadly construed. We aim to highlight new and emerging research opportunities for the machine learning community that arise from the evolving needs for personalization.

The driving factor for new opportunities in personalization is the rapid growth and sophistication of online digital systems that users can interact with (and the resulting interaction data). Personalization first gained significant traction as a way to improve the quality of information retrieval and recommender systems. As the diversity of online content has grown, the development of more effective personalized retrieval and recommender systems remains an important goal. In addition, the emergence of new types of digital systems has expanded the opportunities for personalization to be applied to a wider range of interaction paradigms. Examples of new paradigms include data organization services such as CiteULike and Pinterest, online tutoring systems, and question & answer services such as Quora.

Because the primary asset that enables personalization is the wealth of interaction data, machine learning will play a central role in virtually all future research directions. As a premier machine learning conference, NIPS is an ideal venue for hosting this workshop. Interaction data can pose many interesting machine learning challenges, such as the sheer scale, the multi-task nature of personalizing to populations of users, the exploration/exploitation trade-off when personalizing “on-the-fly”, structured prediction such as formulating a lesson plan in tutoring systems, how to interpret implicit feedback for unbiased learning from interaction data, and how to infer complex sensemaking goals from observing fine-grained interaction sequences.

In summary, our technical topics of interest include (but are not limited to):
- Learning fine-grained representations of user preferences
- Large-scale personalization
- Interpreting observable human behavior
- Interactive algorithms for “on-the-fly” personalization
- Learning to personalize using rich user interactions
- Modeling complex sensemaking goals
- Applications beyond conventional recommender systems

WEBSITE: https://sites.google.com/site/nips2014personalization

Perturbations, Optimization, and Statistics

ORGANIZERS:

Tamir Hazan — University of Haifa
George Papandreou — Toyota Technological Institute
Daniel Tarlow — Microsoft Research

LOCATION: Level 5; room 515b,c

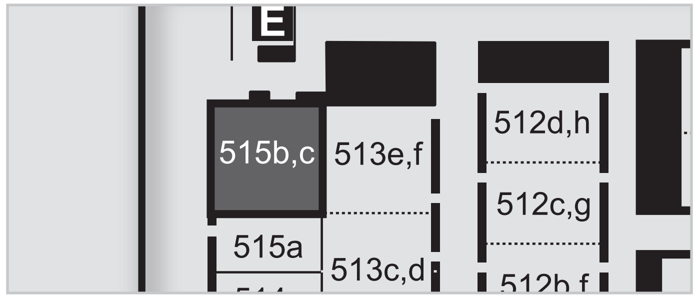

ABSTRACT:

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.

In the last couple of years, at the highly successful NIPS workshops 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 previous workshops 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 look at exciting new developments related to the above core themes.

More generally, we shall specifically be interested in understanding the following issues:

* Modeling: which models lend efficient learning by perturbations?
* Regularization: whether randomness can be replaced by other mathematical object while keeping the computational and statistical guarantees?
* Robust optimization: how stochastic and adversarial perturbations affect the learning outcome?
* Dropout: How stochastic dropout regularizes online learning tasks?
* Sampling: how perturbation can be applied to sample from continuous spaces?

WEBSITE: http://www.cs.toronto.edu/~dtarlow/pos14

SATURDAY WORKSHOPS

SATURDAY WORKSHOP
8:30AM - 6:30PM

3rd NIPS Workshop on Probabilistic Programming

ORGANIZERS:

Vikash Mansinghka	MIT
Josh Tenenbaum	MIT
Dan Roy	University of Toronto
Thomas Dietterich	Oregon State University
Stuart Russell	UC Berkeley

LOCATION: Level 5; room 512 d,h

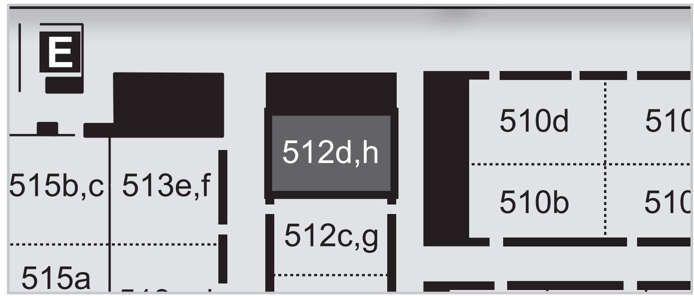

ABSTRACT:

Probabilistic models and approximate inference algorithms are powerful, widely-used tools, central to fields ranging from machine learning to robotics to genetics. However, even simple variations on models and algorithms from the standard machine learning toolkit can be difficult and time-consuming to design, specify, analyze, implement, optimize and debug. The emerging field of probabilistic programming aims to address these challenges by developing formal languages and software systems that integrate key ideas from probabilistic modeling and inference with programming languages and Turing-universal computation.

Over the past two years, the field has rapidly grown and begun to mature. Systems have developed enough that they are seeing significant adoption in real-world applications while also highlighting the need for research in profiling, testing, verification and debugging. New academic and industrial languages have been developed, yielding new applications as well as new technical problems. General-purpose probabilistic programming languages have emerged, complementing previous work focused on specific domains; some offer programmable inference so that experts can take over where automatic techniques fail. The field is has also begun to explore new AI architectures that perform probabilistic reasoning or hierarchical Bayesian approaches for inductive learning over rich data structures and software simulators.

This workshop will survey recent progress, emphasizing results from the ongoing DARPA PPAML program on probabilistic programming. A key theme will be articulating formal connections between probabilistic programming and other fields central to the NIPS community.

WEBSITE: http://probabilistic-programming.org/wiki/NIPS*2014_Workshop

SATURDAY WORKSHOP
8:30AM - 6:30PM

4th Workshop on Automated Knowledge Base Construction (AKBC)

ORGANIZERS:

Partha Talukdar	Indian Institute of Science
Sebastian Riedel	University College London
Sameer Singh	University of Washington
Fabian Suchanek	Paris-Saclay University

ABSTRACT:

Extracting knowledge from Web pages, and integrating it into a coherent knowledge base (KB) is a task that spans the areas of natural language processing, information extraction, information integration, databases, search, and machine learning. Recent years have seen significant advances here, both in academia and in the industry. Most prominently, all major search engine providers (Yahoo!, Microsoft Bing, and Google) nowadays experiment with semantic KBs. Our workshop serves as a forum for researchers on knowledge base construction in both academia and industry.

Unlike many other workshops, our workshop puts less emphasis on conventional paper submissions and presentations, but more on visionary papers and discussions. In addition, one of its unique characteristics is that it is centered on keynotes by high-profile speakers. AKBC 2010, AKBC 2012, and AKBC 2013 each had a dozen invited talks from leaders in this area from academia, industry, and government agencies. We had senior invited speakers from Google, Microsoft, Yahoo, several leading universities (MIT, University of Washington, CMU, University of Massachusetts, and more), and DARPA. With this year's proposal, we would like to resume this positive experience. By inviting established researchers for keynotes, and by focusing particularly on vision paper submissions, we aim to provide a vivid forum of discussion about the field of automated knowledge base construction.

LOCATION: Level 5; room 515

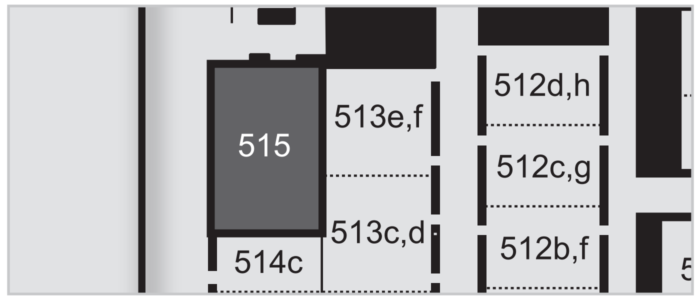

Topics of interest:

* Machine learning on text; unsupervised, lightly-supervised and distantly-supervised learning; learning from naturally-available data.
* Human-computer collaboration in knowledge base construction; automated population of wikis.
* Inference for graphical models and structured prediction; scalable approximate inference.
* Information extraction; open information extraction, named entity extraction; ontology construction;
* Entity resolution, relation extraction, information integration; schema alignment; ontology alignment; monolingual alignment, alignment between knowledge bases and text;
* Pattern analysis, semantic analysis of natural language, reading the web, learning by reading databases; distributed information systems; probabilistic databases;
* Scalable computation; distributed computation.
* Queries on mixtures of structured and unstructured data; querying under uncertainty;
* Dynamic data, online/on-the-fly adaptation of knowledge.
* Languages, toolkits and systems for automated knowledge base construction.
* Demonstrations of existing automatically-built knowledge bases.

WEBSITE: http://akbc.ws

SATURDAY WORKSHOP
8:30AM - 6:30PM

Advances in Variational Inference

ORGANIZERS:

Shakir Mohamed	UC Berkeley
Charles Blundell	UC Berkeley
Google DeepMind	UC Berkeley
Tamara Broderick	UC Berkeley
Michael Jordan	UC Berkeley
Matthew Hoffman	Adobe Systems
David Blei	Columbia University

LOCATION: Level 5; room 510a

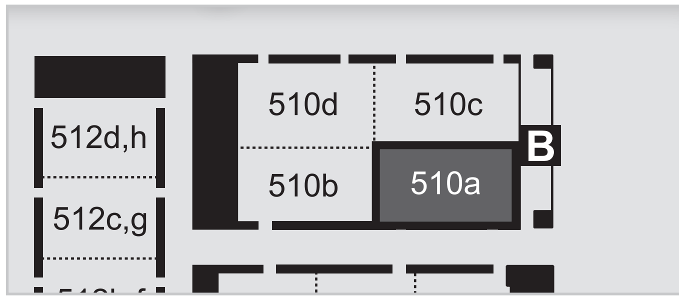

ABSTRACT:

The ever-increasing size of data sets has resulted in an immense effort in machine learning and statistics to develop more powerful and scalable probabilistic models. Efficient inference remains a challenge and limits the use of these models in large-scale scientific and industrial applications. Traditional unbiased inference schemes such as Markov chain Monte Carlo (MCMC) are often slow to run and difficult to evaluate in finite time. In contrast, variational inference allows for competitive run times and more reliable convergence diagnostics on large-scale and streaming data—while continuing to allow for complex, hierarchical modelling. This workshop aims to bring together researchers and practitioners addressing problems of scalable approximate inference to discuss recent advances in variational inference, and to debate the roadmap towards further improvements and wider adoption of variational methods.

The recent resurgence of interest in variational methods includes new methods for scalability using stochastic gradient methods, extensions to the streaming variational setting, improved local variational methods, inference in non-linear dynamical systems, principled regularisation in deep neural networks, and inference-based decision making in reinforcement learning, amongst others. Variational methods have clearly emerged as a preferred way to allow for tractable Bayesian inference. Despite this interest, there remain significant trade-offs in speed, accuracy, simplicity, applicability, and learned model complexity between variational inference and other approximative schemes such as MCMC and point estimation. In this workshop, we will discuss how to rigorously characterise these tradeoffs, as well as how they might be made more favourable. Moreover, we will address other issues of adoption in scientific communities that could benefit from the use of variational inference including, but not limited to, the development of relevant software packages.

WEBSITE: http://www.variationalinference.org

SATURDAY WORKSHOP
8:30AM - 6:30PM

Analysis of Rank Data: Confluence of Social Choice, Operations Research, and Machine Learning

ORGANIZERS:

Shivani Agarwal	Indian Institute of Science
Arun Rajkumar	Indian Institute of Science
Hossein Azari Soufiani	Harvard University
David Parkes	Harvard University
Guy Bresler	MIT
Devavrat Shah	MIT
Sewoong Oh	UIUC

LOCATION: Level 5; room 513c,d

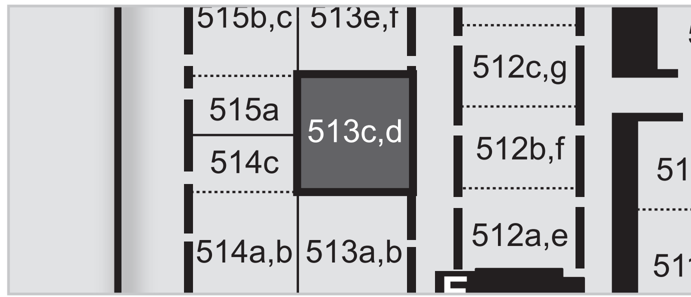

ABSTRACT:

The mathematical analysis and understanding of rank data has been a fascinating topic for centuries, and has been investigated in disciplines as wide-ranging as social choice/voting theory, decision theory, probability, statistics, and combinatorics. In modern times, huge amounts of data are generated in the form of rankings on a daily basis: restaurant ratings, product ratings/comparisons, employer ratings, hospital rankings, doctor rankings, and an endless variety of rankings from committee deliberations (including, for example, deliberations of conference program committees such as NIPS!). These applications have led to several new trends and challenges: for example, one must frequently deal with very large numbers of candidates/alternatives to be ranked, with partial or missing ranking information, with noisy ranking information, with the need to ensure reliability and/or privacy of the rank data provided, and so on.

Given the increasing universality of settings involving large amounts of rank data and associated challenges as above, powerful computational frameworks and tools for addressing such challenges have emerged over the last few years in a variety of areas, including in particular in machine learning, operations research, and computational social choice. Despite the fact that many important practical problems in each area could benefit from the algorithmic solutions and analysis techniques developed in other areas, there has been limited interaction between these areas. Given both the increasing maturity of the research into ranking in these respective areas and the increasing range of practical ranking problems in need of better solutions, it is the aim of this workshop to bring together recent advances in analyzing rank data in machine learning, operations research, and computational social choice under one umbrella, to enable greater interaction and cross-fertilization of ideas.

A primary goal will be to discover connections between recent approaches developed for analyzing rank data in each of the three areas above. To this end, we will have invited talks by leading experts in the analysis of rank data in each area. In addition, we will include perspectives from practitioners who work with rank data in various applied domains on both the benefits and limitations of currently available solutions to the problems they encounter. In the end, we hope to both develop a shared language for the analysis and understanding of rank data in modern times, and identify important challenges that persist and could benefit from a shared understanding.

The topics of interest include:
- discrete choice modeling and revenue management
- voting and social decision making, preference elicitation
- social choice (rank aggregation) versus individual choice (recommendation systems)
- stochastic versus active sampling of preferences
- statistical/learning-theoretic guarantees
- effects of computational approximations

WEBSITE: http://events.csa.iisc.ernet.in/NIPS-14-rankingsws/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Discrete Optimization in Machine Learning

ORGANIZERS:

Stefanie Jegelka	UC Berkeley
Andreas Krause	ETH Zurich
Jeff Bilmes	University of Washington, Seattle

LOCATION: Level 5; room 514

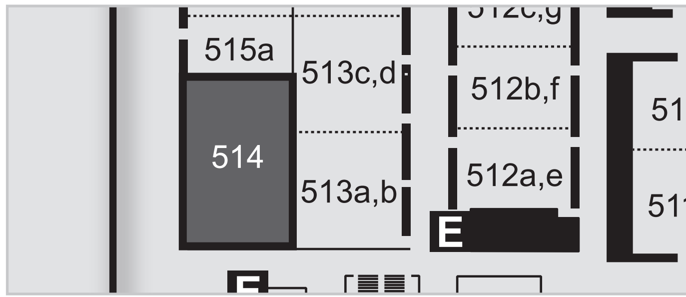

ABSTRACT:

We propose to hold a workshop that will bring together scientists interested in both discrete/combinatorial optimization and machine learning, and in doing so bringing together 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. The workshop aims to ask questions that enable communication across such fields. In particular, this year we aim to address how the investigation of combinatorial structures allows to capture complex, high-order dependencies in discrete learning problems, both in theory and applications.

WEBSITE: http://discml.cc/

SATURDAY WORKSHOP
8:30AM - 6:30PM

High-energy particle physics, machine learning, and the HiggsML data challenge (HEPML)

ORGANIZERS:

Balázs Kégl	University of Paris-Sud/CNRS
Glen Cowan	University of London
Royal Holloway	University of London
Cecile Germain	Université Paris Sud
Isabelle Guyon	Clopinet
David Rousseau	CNRS

LOCATION: Level 5; room 511c

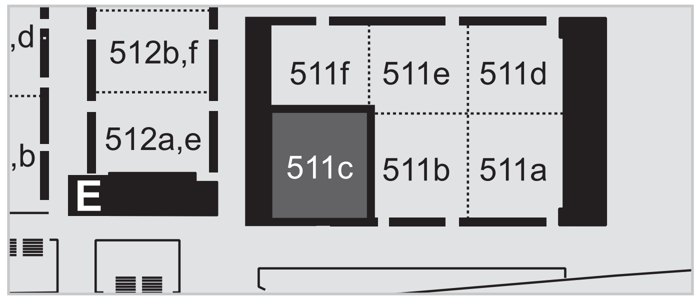

ABSTRACT:

What is high-energy particle physics

The objective of high-energy particle physics (HEP) is to study the basic constituents of matter, largely within the theory called the Standard Model and its possible extensions. The main experimental tools are particle accelerators and colliders in which beams of particles are accelerated to very high kinetic energy and collided into other particles. The particles resulting from the collision are then detected in particle detectors consisting mainly of track detectors (high-resolution devices in which the paths of individual charged particles can be separated) and calorimeters (measuring the energy of particles or groups of particles). From these raw measurements, different events (mainly particle decays and collisions) are reconstructed, the whole “picture” is compared to model predictions, and model parameters (for example, the existence and the mass of new particles) are inferred from comparing a large statistics of collision events to simulated events.

Machine learning and high-energy particle physics

Particle physics is one of the posterboys of today’s data science revolution. Indeed, large-scale HEP experiments assimilated computational paradigms a long time ago: both simulators and semi-automatic data analysis techniques have been applied widely for decades. In particular, nonparametric classification and regression are now routinely used as parts of the reconstruction (inference) chain. More recently, state-of- the-art budgeted learning techniques have also started to be used for real-time event selection, and this year we have also seen the first use of deep learning techniques in particle searches. Nevertheless, most of these applications went largely unnoticed by the machine learning (ML) community. The goal of this workshop is to contribute to improving the dialog between these two communities. The expected outcome is twofold: we hope to introduce a set of exciting scientific questions and open problems in computational statistics and machine learning to the NIPS community, and, at the same time, to transmit some of the state-of-the-art techniques developed in ML to the HEP community.

The HiggsML challenge

https://www.kaggle.com/c/higgs-boson - The main motivation to propose this workshop now is the HiggsML challenge, organized by the same group as the WS organizers. The challenge, hosted by Kaggle, has drawn a remarkably large audience (with 1500+ teams it is one of the all-time most popular Kaggle challenges) and large coverage both in the social and in the classical media. The goal of the challenge is to improve the procedure that classifies events produced by the decay of the Higgs boson versus events produced by other (background) processes, based on a training set of 250000 examples. The challenge is a premier: it is the first time that a CERN experiment (ATLAS) made public such a large set of the official event and detector simulations. It also features a unique formal objective representing an approximation of the median significance (AMS) of a discovery (counting) test, which generates interesting algorithmic/theoretical questions beyond the usual challenges of finding and tuning the best classification algorithm.

WEBSITE: https://sites.google.com/site/hepml14/home

SATURDAY WORKSHOP
8:30AM - 6:30PM

Human Propelled Machine Learning

ORGANIZERS:

Richard Baraniuk	Rice University
Divyanshu Vats	Rice University
Andrew Waters	Rice University
Andrew Lan	Rice University
Michael Mozer	University of Colorado
Christoph Studer	Cornell University

LOCATION: Level 5; room 511f

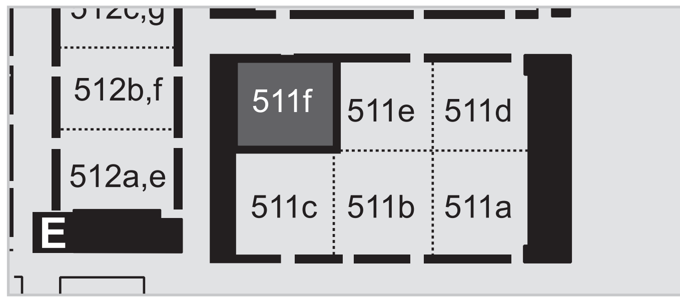

ABSTRACT:

In typical applications of machine learning (ML), humans typically enter the process at an early stage, in determining an initial representation of the problem and in preparing the data, and at a late stage, in interpreting and making decisions based on the results. Consequently, the bulk of the ML literature deals with such situations. Much less research has been devoted to ML involving "humans-in-the-loop," where humans play a more intrinsic role in the process, interacting with the ML system to iterate towards a solution to which both humans and machines have contributed. In these situations, the goal is to optimize some quantity that can be obtained only by evaluating human responses and judgments. Examples of this hybrid, "human-in-the-loop" ML approach include:

- ML-based education, where a scheduling system acquires information about learners with the goal of selecting and recommending optimal lessons;
- Adaptive testing in psychological surveys, educational assessments, and recommender systems, where the system acquires testees' responses and selects the next item in an adaptive and automated manner;
- Interactive topic modeling, where human interpretations of the topics are used to iteratively refine an estimated model;
- Image classification, where human judgments can be leveraged to improve the quality and information content of image features or classifiers.

The key difference between typical ML problems and problems involving "humans-in-the-loop" and is that in the latter case we aim to fit a model of human behavior as we collect data from subjects and adapt the experiments we conduct based on our model fit. This difference demands flexible and robust algorithms and systems, since the resulting adaptive experimental design depends on potentially unreliable human feedback (e.g., humans might game the system, make mistakes, or act lazily or adversarially). Moreover, the "humans-in-the-loop" paradigm requires a statistical model for human interactions with the environment, which controls how the experimental design adapts to human feedback; such designs are, in general, difficult to construct due to the complex nature of human behavior. Suitable algorithms also need to be very accurate and reliable, since humans prefer a minimal amount of interaction with ML systems; this aspect also prevents the use of computationally intensive parameter selection methods (e.g., a simple grid search over the parameter space). These requirements and real-world constraints render "humans-in-the-loop" ML problems much more challenging than more standard ML problems.

In this workshop, we will focus on the emerging new theories, algorithms, and applications of human-in-the-loop ML algorithms. Creating and estimating statistical models of human behavior and developing computationally efficient and accurate methods will be a focal point of the workshop. This human-behavior aspect of ML has not been well studied in other fields that rely on human inputs such as active learning and experimental design. We will also explore other potential interesting applications involving humans in the loop in different fields, including, for example, education, crowdsourcing, mobile health, pain management, security, defense, psychology, game theory, and economics.

The goal of this workshop is to bring together experts from different fields of ML, cognitive and behavioral sciences, and human-computer interaction (HCI) to explore the interdisciplinary nature of research on this topic. In particular, we aim to elicit new connections among these diverse fields, identify novel tools and models that can be transferred from one to the other, and explore novel ML applications that will benefit from the human-in-the-loop of ML algorithms paradigm. We believe that a successful workshop will lead to new research directions in a variety of areas and will also inspire the development of novel theories and tools.

WEBSITE: http://dsp.rice.edu/HumanPropelledML_NIPS2014

SATURDAY WORKSHOP
8:30AM - 6:30PM

Large-scale Reinforcement Learning and Markov Decision Problems

ORGANIZERS:

Yasin Abbasi — Queensland University
Peter Bartlett — UC Berkeley
Mohammad Ghavamzadeh — Adobe Research & INRIA
Ambuj Tewari — University of Michigan
Benjamin Van Roy — Stanford University

LOCATION: Level 5; room 511d

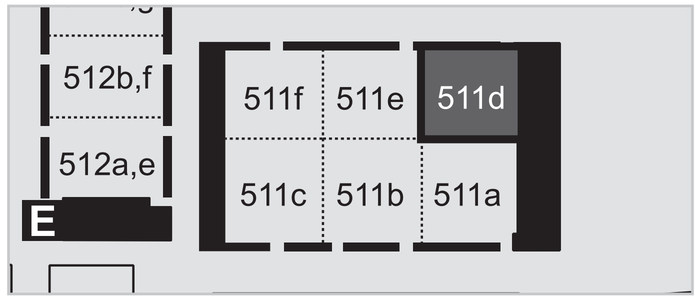

ABSTRACT:

Reinforcement learning (RL) and MDPs have been topics of intense research since the middle of the last century. It was shown that Dynamic Programming (DP) [B, H] gives the optimal policy and its computational cost is polynomial in the number of states and actions. This polynomial dependence on the size of the state space limits exact DP to small state problems. Modern applications of RL need to deal with large state problems that arise in many areas ranging from robotics to medical trials to finance.

Solving a large state MDP problem can be computationally intractable in the worst case [PT, CT]. Despite these negative results, several algorithms are shown to perform remarkably well in certain large state problems. Examples are UCT algorithm of Kocsis and Szepesvari [KS] applied in heuristic search and games, Rapidly exploring Random Trees (RRT) of LaValle and Kuffner [LK] in motion planning, policy gradient methods applied in robotics [KP, GE], approximate linear programming (ALP) applied in queuing networks [FV, DFM], and approximate dynamic programming applied in very large scale industrial applications [P]. These algorithms are developed mostly independently in different communities. Despite some similarities, the relation between them and what makes them effective is not very clear.

An important feature of these methods is that they take an alternative approach than the typically used approaches based on the notion of optimality. Instead, they search for a reasonably good policy in a smaller space. For instance, ALP and ADP methods approximate the value function by a linear combination of a number of features [S, SSM, FV]. More recently [ABM] propose to search for the best policy in a low-dimensional policy space. As another example, RRTs try to find only a feasible path (a path that connects the starting state to the goal state), rather than the optimal (shortest) path in the search space. It turns out that finding a feasible path can be substantially easier than finding the optimal path in certain classes of problems. The research on MDP problems is conducted in different communities: machine learning, operations research/optimization, robotics, and heuristic search/planning.

The computational problem that we discussed above is the focus of optimization and ADP communities. The challenge is to find a computationally efficient way to achieve something that would be easy if we had infinite computational resources. In reinforcement learning, we encounter additional statistical challenges; even if we have infinite computational power, it is not clear how we should best make inferences from observations and select actions that balance between exploration and exploitation.

This workshop will bring researchers from different communities together to discuss and exchange ideas about effective approaches and open problems in large scale MDP problems. Further, this multidisciplinary gathering will provide a perfect opportunity to clarify the confusion in the community about the difference and relation between RL and ADP.

WEBSITE: https://sites.google.com/site/rlmdp2014/

Machine Learning in Computational Biology

ORGANIZERS:

Anna Goldenberg	University of Toronto
Sara Mostafavi	Harvard University
Oliver Stegle	EMBL-EBI
Su-In Lee	University of Washington
Martin Renqiang Min	NEC Labs America

LOCATION: Level 5; room 510b

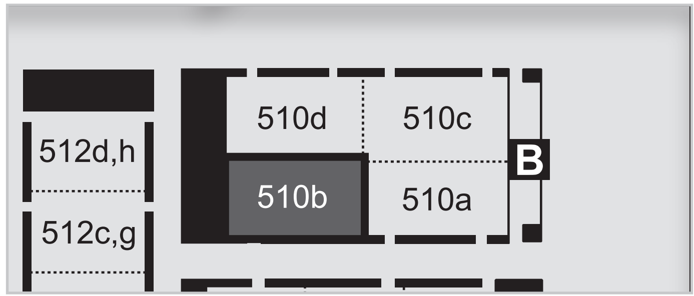

ABSTRACT:

The field of computational biology has seen dramatic growth over the past few years. A wide range of high-throughput technologies developed in the last decade now enable us to measure parts of a biological system at various resolutions—at the genome, epigenome, transcriptome, and proteome levels. These data are high-dimensional, heterogeneous, and are impacted by a range of confounding factors, presenting new challenges for standard learning and inference approaches. Therefore, fully realizing the scientific and clinical potential of these data requires development of novel supervised and unsupervised learning methods that are scalable, can accommodate heterogeneity, are robust to systematic noise and confounding factors, and provide mechanistic insights.

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. We will also have the usual rigorous screening of 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: www.mlcb.org

SATURDAY WORKSHOP
8:30AM - 6:30PM

MLINI 2014 - 4th NIPS Workshop on Machine Learning and Interpretation in Neuroimaging: Beyond the Scanner

ORGANIZERS:

Irina Rish	IBM T.J. Watson Research Center
Georg Langs	Medical University of Vienna
Brian Murphy	Queen’s University Belfast
Guillermo Cecchi	IBM Research
Kai-min Chang	Carnegie Mellon University
Leila Wehbe	Carnegie Mellon University

LOCATION: Level 5; room 513e,f

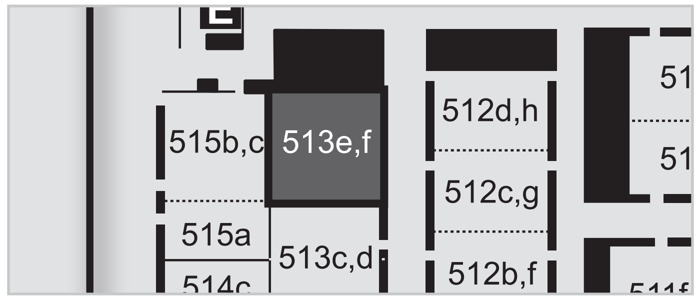

ABSTRACT:

Aim MLINI workshop focuses on machine learning approaches in neuroscience, neuroimaging, with a specific extension to behavioral experiments and psychology. The special topic of this year is “Going Beyond the Scanner”, which includes making inference about the subject’s mental states from ‘’cheap’’ data such as subject’s speech and/or text, audio, video, EEG and other wearable devices.

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. It is also equally important that practical applications in these fields motivate a rapidly evolving line or research in the machine learning community. In parallel, there is an intense interest in learning more about brain function in the context of rich naturalistic environments and scenes. Efforts to go beyond highly specific paradigms that pinpoint a single function, towards schemes for measuring the interaction with natural and more varied scene are made. In this context, many controversies and open questions exist.

The goal of the workshop is to pinpoint the most pressing issues and common challenges across the fields, and to sketch future directions and open questions in the light of novel methodology. The proposed 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 the previous workshops in this series, MLINI ‘11, MLINI’12, and MLINI’13, 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 neuroimaging scope and investigate the role of behavioral models and psychology, including topics such as psycholinguistics.

WEBSITE: https://sites.google.com/site/mlininips2014/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Modern Nonparametrics 3: Automating the Learning Pipeline

ORGANIZERS:

Arthur Gretton	UCL
Samory Kpotufe	TTI Chicago
Andrew Wilson	Carnegie Mellon University
Han Liu	Princeton University
Mladen Kolar	University of Chicago
Zoltan Szabo	University College London
Eric Xing	CMU

LOCATION: Level 5; room 511e

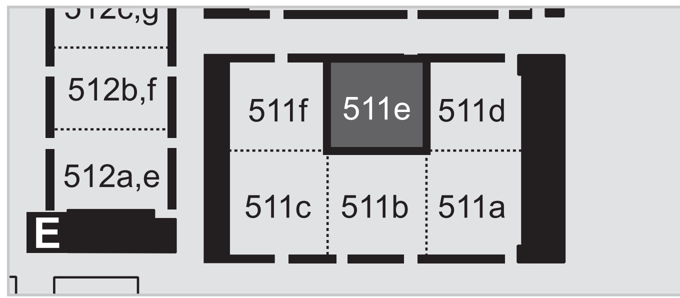

ABSTRACT:

Nonparametric methods (kernel methods, kNN, classification trees, etc) are designed to handle complex pattern recognition problems. Such complex problems arise in modern applications such as genomic experiments, climate analysis, robotic control, social network analysis, and so forth. In fact, contemporary statistical procedures are making inroads into a variety of modern application areas as part of solutions to larger problems. As such there is a growing need for statistical procedures that can be used “off-the-shelf”, i.e. procedures with as few parameters as possible, or better yet, procedures which can “self-tune” to a particular application at hand.

The problem of devising ‘parameter-free’ procedures has been addressed in separate areas of the pattern-recognition literature under various names and different emphasis. In traditional statistics, much effort has gone into so called “adaptive” procedures which can attain optimal risks over large sets of models of increasing complexity. Examples are model selection approaches based on penalized empirical risk minimization, approaches based on stability of estimates (e.g. Lepski’s methods), thresholding approaches under sparsity assumptions, and model averaging approaches. Most of these approaches rely on having tight bounds on the risk of learning procedures (under any parameter setting), hence other approaches concentrate on tight estimations of the actual risks, e.g., Stein’s risk estimators, bootstrapping methods, data dependent learning bounds.

In theoretical machine learning, much of the work has focused on proper tuning of the actual optimization procedures used to minimize (penalized) empirical risks. In particular, great effort has gone into the automatic setting of important tuning parameters such as ‘learning rates’ and ‘step sizes’.

Another approach out of machine learning arises in the kernel literature for ‘automatic representation learning’. The aim of the approach, similar to theoretical work on model selection, is to automatically learn an appropriate (kernel) transformation of the data for use with kernel methods such as SVMs or Gaussian processes.

In practice, the simplest self-tuning procedures take the form of cross-validation and variants. Cross-validation can however be expensive in practice, and impractical in various constrained settings -- e.g., streaming settings, in settings with large amounts of tuning parameters, and generally in unsupervised learning problems. More generally, many existing self-tuning or parameter-free methods are unfortunately expensive given large modern data sizes and dimensionality, while the cheaper methods tend to self-tune only to small model classes. Ideally we would want self-tuning procedures that can adapt to easy or difficult (nonparametric) problems, while satisfying the practical constraints of modern applications.

A main aim of this workshop is to cover the various approaches proposed so far towards automating the learning pipeline, and the practicality of these approaches in light of modern constraints. We are particularly interested in understanding whether large datasizes and dimensionality might help the automation effort since such datasets in fact provide more information on the patterns being learned.

Through a number of invited and contributed talks and a focused panel discussion, we plan to bring together both theoretical and applied researchers to discuss these challenges in detail, share insight on existing solutions, and lay out some of the important future directions towards answering the demands of modern applications.

WEBSITE: https://sites.google.com/site/nips2014modernnonparametric/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Networks: From Graphs to Rich Data

ORGANIZERS:

David Choi	Carnegie Mellon University
Aaron Clauset	University of Colorado
Leto Peel	University of Colorado
Edo Airoldi	Harvard University
Johan Ugander	Microsoft Research

LOCATION: Level 5; room 510c

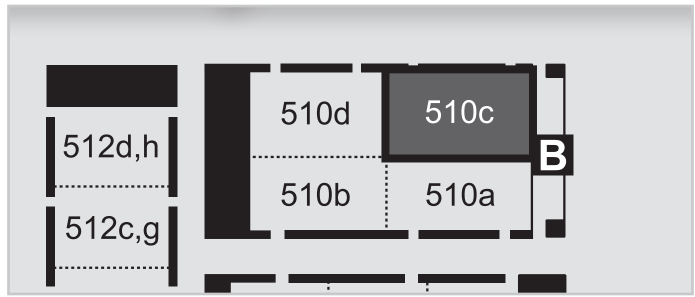

ABSTRACT:

Problems involving networks and massive network datasets motivate some of the most difficult and exciting inferential challenges today, from social, economic, and biological domains. Modern network data are often more than just vertices and edges, containing rich information on vertex attributes, edge weights, and characteristics that change over time. Enormous in size, detail, and heterogeneity, these networks 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 network analysis is a young and highly cross-disciplinary field.

This workshop aims to bring together a diverse and cross-disciplinary set of researchers to discuss recent advances and 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 informational systems. We also welcome empirical studies in applied domains such as the social sciences, biology, medicine, neuroscience, physics, finance, social media, and economics. And, we are particularly interested in research that unifies the study of both structure and content in rich network datasets.

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.

The workshop's primary goal is to facilitate the technical maturation of network analysis, promote greater technical sophistication and practical relevance, and identify future directions of research. This workshop will thus 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 in goals.

The NIPS community is well positioned as a middle ground for effective dialog between applied and methodological concerns. We aim to further leverage this position to facilitate an open, cross-disciplinary discussion among researchers to assess progress and stimulate further debate on networks. We believe these efforts will ultimately yield novel modeling approaches and the identification of new applications or open problems that will guide future research in networks.

WEBSITE: http://tinyurl.com/NetworksNIPS2014

NIPS'14 Workshop on Crowdsourcing and Machine Learning

ORGANIZERS:

Adish Singla	ETH, Zurich
Chien-Ju Ho	UCLA/Harvard
David Parkes	Harvard University
Nihar Shah	UC Berkeley
Dengyong Zhou	Microsoft Research

LOCATION: Level 5; room 511a

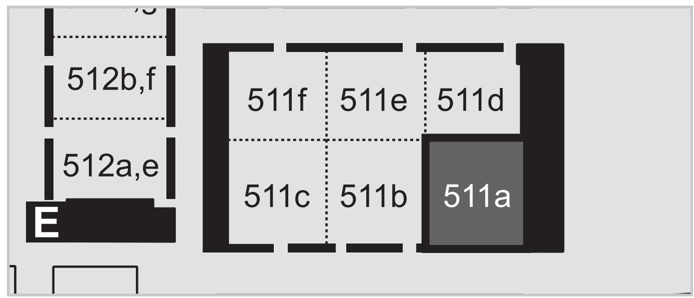

ABSTRACT:

Crowdsourcing aims to combine human knowledge and expertise with computing to help solve problems and scientific challenges that neither machines nor humans can solve alone. In addition to a number of human-powered scientific projects, including GalaxyZoo, eBird, and Foldit, crowdsourcing is impacting the ability of academic researchers to build new systems and run new experiments involving people, and is also gaining a lot of use within industry for collecting training data for the purpose of machine learning. There are a number of online marketplaces for crowdsourcing, including Amazon's Mechanical Turk, ODesk and MobileWorks. The fundamental question that we plan to explore in this workshop is: **How can we build systems that combine the intelligence of humans and the computing power of machines for solving challenging scientific and engineering problems?**

The goal is to improve the performance of complex human-powered systems by making them more efficient, robust, and scalable. Current research in crowdsourcing often focuses on micro-tasking (for example, labeling a set of images), designing algorithms for solving optimization problems from the job requester's perspective and with simple models of worker behavior. However, the participants are people with rich capabilities including learning, collaboration and so forth, suggesting the need for more nuanced approaches that place special emphasis on the participants. Such human-powered systems could involve large numbers of people with varying expertise, skills, interests, and incentives. This poses many interesting research questions and exciting opportunities for the machine learning community. The goal of this workshop is to foster these ideas and work towards this goal by bringing together experts from the field of machine learning, cognitive science, economics, game theory, and human-computer interaction.

Topics of interests in the workshop include:

- **Social aspects and collaboration:** How can systems exploit the social ties of the underlying participants or users to create incentives for users to collaborate? How can online social networks be used to create tasks with a gamification component and engage users in useful activities? With ever-increasing time on the Internet being spent on online social networks, there is a huge opportunity to elicit useful contributions from users at scale, by carefully designing tasks.

- **Incentives, pricing mechanisms and budget allocation**: How to design the right incentive structure and pricing policies for participants that maximize the satisfaction of participants as well as utility of the job requester for a given budget? How can techniques from machine learning, economics and game theory be used to learn optimal pricing policies and to infer optimal incentive designs?

- **Learning by participants:** How can we use insights from machine learning to build tools for training and teaching the participants for carrying out complex or difficult tasks? How can this training be actively adapted based on the skills or expertise of the participants and by tracking the learning process?

- **Peer prediction and knowledge aggregation:** How can complex crowdsourcing tasks be decomposed into simpler micro-tasks? How can techniques of peer prediction be used to elicit informative responses from participants and incentivize effort? Can we design models and algorithms to effectively aggregate responses and knowledge, especially for complex tasks?

- **Privacy aspects:** The question of privacy in human-powered systems has often been ignored and we seek to understand the privacy aspects both from job requester as well as privacy of the participants. How can a job requester (such as firm interested in translating legal documents) carry out crowdsourcing tasks without revealing private information to the crowd? How can systems negotiate the access to private information of participants (such as the GPS location in community sensing applications) in return of appropriate incentives?

- **Open theoretical questions and novel applications:** What are the open research questions, emerging trends and novel applications related to design of incentives in human computation and crowdsourcing systems?

We expect diverse participation from researchers with a wide variety of scientific interests spanning economics, game theory, cognitive science, and human-computer interaction. Given the widespread use of crowdsourcing in the industry, such Amazon, Google and Bing, we expect active participation from industry.

WEBSITE: http://crowdwisdom.cc/nips2014/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Novel Trends and Applications in Reinforcement Learning

ORGANIZERS:

Csaba Szepesvari	University of Alberta
Marc Deisenroth	Imperial College London
Sergey Levine	UC Berkeley
Pedro Ortega	University of Pennsylvania

LOCATION: Level 5; room 512 a,e

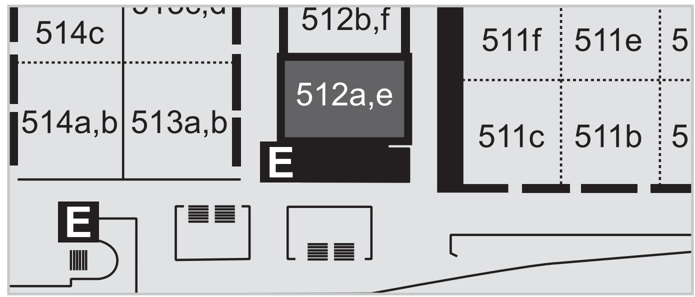

ABSTRACT:

The last decade has witnessed a series of technological advances: social networks, cloud servers, personalized advertising, autonomous cars, personalized healthcare, robotics, security systems, just to name a few. These new technologies have in turn substantially reshaped our demands from adaptive reinforcement learning systems, defining novel yet urgent challenges. In response, a wealth of novel ideas and trends have emerged, tackling problems such as modelling rich and high-dimensional dynamics, life-long learning, resource-bounded planning, and multi-agent cooperation.

The objective of the workshop is to provide a platform for researchers from various areas (e.g., deep learning, game theory, robotics, computational neuroscience, information theory, Bayesian modelling) to disseminate and exchange ideas, evaluating their advantages and caveats. In particular, we will ask participants to address the following questions:

1) What is the future of reinforcement learning?
2) What are the most important challenges?
3) What tools do we need the most?

A final panel discussion will then review the provided answers and focus on elaborating a list of trends and future challenges. Recent advances will be presented in short talks and a poster session based on contributed material.

WEBSITE: http://tcrl14.wordpress.com/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Optimal Transport and Machine Learning

ORGANIZERS:

Marco Cuturi — Kyoto University
Gabriel Peyré — Université Paris Dauphine
Justin Solomon — Stanford University

LOCATION: Level 5; room 512 c,g

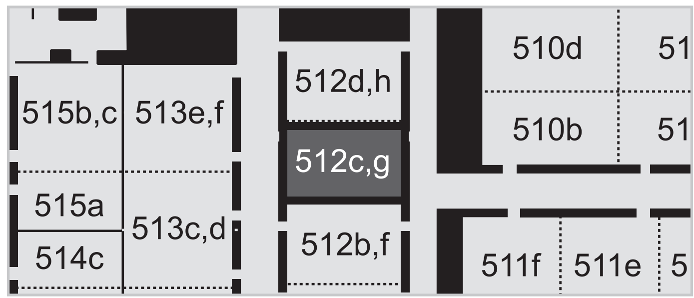

ABSTRACT:

Optimal transport (OT) has emerged as a novel tool to solve problems in machine learning and related fields, e.g. graphics, statistics, data analysis, computer vision, economics and imaging.

In particular, the toolbox of OT (including for instance the Wasserstein/Earth Mover's Distances) offers robust mathematical techniques to study probability measures and compare complex objects described using bags-of-features representations.

Scaling OT algorithms to datasets of large dimension and sample size presents, however, a considerable computational challenge. Taking for granted that these challenges are partially solved, there remains many salient open research questions on how to integrate OT in statistical methodologies (dimensionality reduction, inference, modeling) beyond its classical use in retrieval. OTML 2014 will be the first international workshop to address state-of-the-art research in this exciting area.

WEBSITE: http://www.iip.ist.i.kyoto-u.ac.jp/OTML2014/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Representation and Learning Methods for Complex Outputs

ORGANIZERS:

Yuhong Guo	Temple University
Dale Schuurmans	University of Alberta
Richard Zemel	University of Toronto
Kilian Weinberger	Washington University, St. Louis

LOCATION: Level 5; room 512 b,f

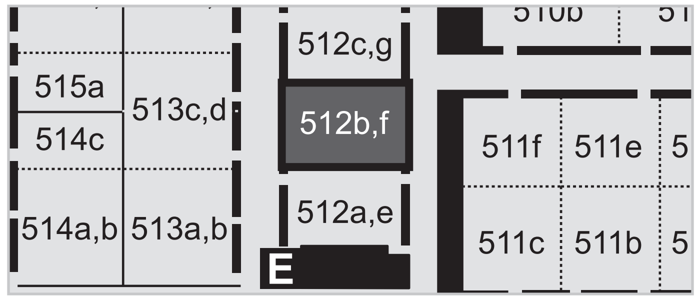

ABSTRACT:

Learning problems that involve complex outputs are becoming increasingly prevalent in machine learning research. For example, work on image and document tagging now considers thousands of labels chosen from an open vocabulary, with only partially labeled instances available for training. Given limited labeled data, these settings also create zero-shot learning problems with respect to omitted tags, leading to the challenge of inducing semantic label representations. Furthermore, prediction targets are often abstractions that are difficult to predict from raw input data, but can be better predicted from learned latent representations. Finally, when labels exhibit complex inter-relationships it is imperative to capture latent label relatedness to improve generalization.

This workshop will bring together separate communities that have been working on novel representation and learning methods for problems with complex outputs. Although representation learning has already achieved state of the art results in standard settings, recent research has begun to explore the use of learned representations in more complex scenarios, such as structured output prediction, multiple modality co-embedding, multi-label prediction, and zero shot learning. Unfortunately, these emerging research topics have been conducted in separate sub-areas, without proper connections drawn to similar ideas in other areas, hence general methods and understanding have not yet emerged from the disconnected pursuits. The aim of this workshop is to identify fundamental strategies, highlight differences, and identify the prospects for developing a set of systematic theory and methods for learning problems with complex outputs. 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.

Relevant topics include:
- Multi-label learning with large and/or incomplete output spaces
- Zero-shot learning
- Co-embedding
- Learning output kernels
- Output structure learning

WEBSITE: https://sites.google.com/site/complexoutputs2014/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Riemannian Geometry in Machine Learning, Statistics and Computer Vision

ORGANIZERS:

Minh Ha Quang	Istituto Italiano di Tecnologia
Vikas Sindhwani	IBM Research
Vittorio Murino	Italian Institute of Technology

LOCATION: Level 5; room 511b

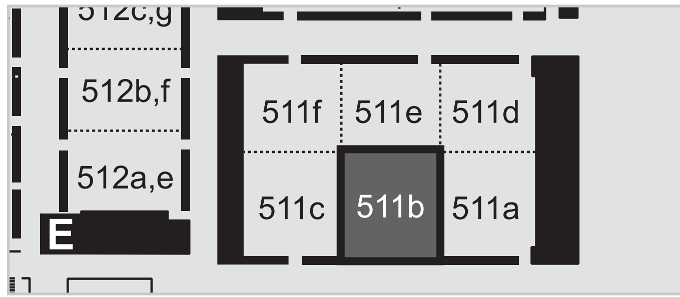

ABSTRACT:

Traditional machine learning and data analysis methods often assume that the input data can be represented by vectors in Euclidean space. While this assumption has worked well for many applications, researchers have increasingly realized that if the data is intrinsically non-Euclidean, ignoring this geometrical structure can lead to suboptimal results.

In the existing literature, there are two common approaches for exploiting data geometry when the data is assumed to lie on a Riemannian manifold.

In the first direction, often referred to as manifold learning, the data is assumed to lie on an unknown Riemannian manifold and the structure of this manifold is exploited through the training data, either labeled or unlabeled. Examples of manifold learning techniques include Manifold Regularization via the graph Laplacian, Locally Linear Embedding, and Isometric Mapping.

In the second direction, which is gaining increasing importance and success, the Riemannian manifold representing the input data is assumed to be known explicitly. Some manifolds that have been widely used for data representation are: the manifold of symmetric, positive definite matrices, the Grassmannian manifold of subspaces of a vector space, and the Kendall manifold of shapes. When the manifold is known, the full power of the mathematical theory of Riemannian geometry can be exploited in both the formulation of algorithms as well as their theoretical analysis. Successful applications of these approaches are numerous and range from brain imaging and low rank matrix completion to computer vision tasks such as object detection and tracking.

This workshop focuses on the latter direction. We aim to bring together researchers in statistics, machine learning, computer vision, and other areas, to discuss and exchange current state of the art results , both theoretically and computationally, and identify potential future research directions

WEBSITE: http://www.riemanniangeometry2014.eu/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Second Workshop on Transfer and Multi-Task Learning: Theory meets Practice

ORGANIZERS:

Urun Dogan	Microsoft
Marius Kloft	Humboldt University of Berlin
Andres Munoz	Courant Institute of Mathematical Sciences
Francesco Orabona	Toyota Technological Institute
Tatiana Tommasi	KU Leuven

LOCATION: Level 5; room 510d

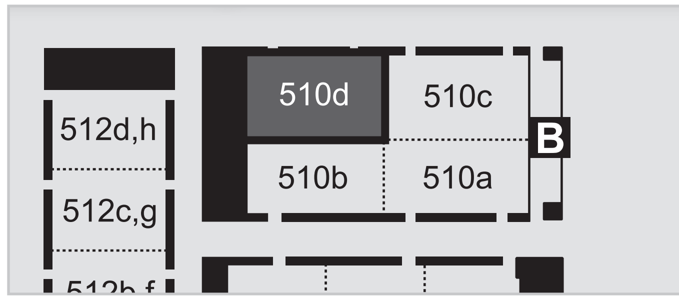

ABSTRACT:

Transfer, domain adaptation and multi-task learning methods have been developed to better exploit the available data at training time, originally moved by the need to deal with a reduced amount of information. Nowadays, gathering data is much easier than in the past thanks to the low price of different acquisition devices (e.g. cameras) and to the World Wide Web that connects million of devices users. Existing methods must embrace new challenges to manage large scale data that do not lack anymore in size but may lack in quality or may continuously change over time. All this comes with several open questions, for instance:

- What are the limits of existing multi-task learning methods when the number of tasks grows while each task is described by only a small bunch of samples (“big T, small n”)?
- Theory vs. practice: can multi-task learning for very big data (n>10^7) be performed with extremely randomized trees?
- What is the right way to leverage over noisy data gathered from the Internet as reference for a new task?
- Can we get an advantage by overcoming the dataset bias and aligning multiple existing but partially related data collections before using them as source knowledge for a new target problem?
- How can an automatic system process a continuous stream of information in time and progressively adapt for life-long learning?
- Since deep learning has demonstrated high performance on large scale data, is it possible to combine it with transfer and multiple kernel learning in a principled manner?
- Can deep learning help to learn the right representation (e.g., task similarity matrix) in kernel-based transfer and 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?
- How can similarities across languages help us adapt to different domains in natural language processing tasks?

After the first workshop edition where we investigated new directions for learning across domains, we want now to call the attention of the machine learning community on the emerging problem of big data and its particular challenges regarding multi-task and transfer learning and its practical effects in many application areas like computer vision, robotics, medicine, bioinformatics etc. where transfer, domain adaptation and multi-task learning have been previously used with success. We will encourage applied researchers to contribute to the workshop in order to create a synergy with theoreticians and lead to a global advancement of the field.

A selection of the papers accepted to the workshop and voted by the reviewers will be re-evaluated also as invited contributions to the planned JMLR special topic on Domain Adaptation, Multi-task and Transfer Learning. The proposal for this special topic is currently under evaluation.

WEBSITE: https://sites.google.com/site/multitaskwsnips2014/

SATURDAY WORKSHOP
8:30AM - 6:30PM

Software Engineering for Machine Learning

ORGANIZERS:

Ryan Turner	Northrop Grumman
Joaquin Quiñonero Candela	Facebook
Xavier Amatriain	Netflix

LOCATION: Level 5; room 513a,b

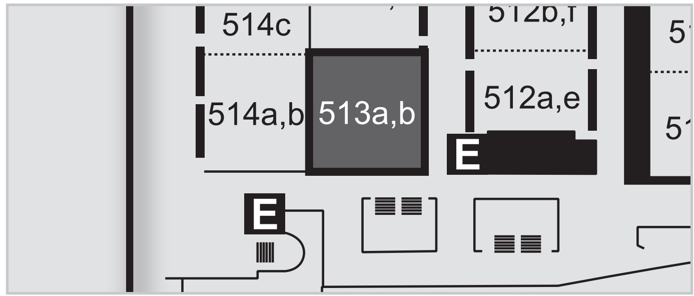

ABSTRACT:

We are organizing a one day NIPS 2014 workshop that will cover topics at the intersection of machine learning and software architecture/ engineering. This intersection is a critical area for deploying machine learning methods in practice, but is often overlooked in the literature. As a result, much of the publicly available code for download is disorganized, undocumented, and buggy. Therefore, it cannot serve as an example of how actual deployed machine-learning-heavy software should be written. Those looking to implement actual software could greatly benefit from a workshop that can provide guidance on software practices.

There are several topics this workshop will cover through contributed and invited talks:

1. Scaling machine learning: Solutions to practical issues involving taking single machine algorithms and making them ready for “big data” by distributing them with Spark or Hadoop/MapReduce are welcome here.
2. Accelerating machine learning prototypes: Methods and tips for moving single machine Matlab/R/Python code to C/C++ code, as well as GPU acceleration.
3. Software paradigms: When is it best to work in an object oriented, procedural, or functional framework when developing machine learning software?
4. When to use probabilistic programming environments? If so, which tool (e.g. Infer.NET, Stan, Church, etc.) is most appropriate for your project requirements?
5. Systematic testing: This is often overlooked but important area for the workshop to cover. Can we develop better methods for systematically testing our methods to make sure they are implemented correctly? This includes unit testing and regression testing. (a) There is a perception among some practitioners that systematic methods like unit tests are not applicable to machine learning because “The whole reason we are doing the computation in the first place is that we do not know the answer.” One goal of this workshop is to try and change that perception with guidance and examples. (b) What are some of the common ways to break down a machine learning project into units where unit testing is possible? Monte Carlo unit tests: Unlike most projects many unit tests in machine learning are Monte Carlo tests. (c) Different inference methods will have their own methods that can be used to test their implementation correctness: VB, EP, MCMC [3; 1], etc.
6. Documentation: How should people in machine learning be doing a better job at documenting their code? Do the usual guidelines for software documentation need to be augmented or modified for machine learning software? Could tools such as literate programming [4] be more useful than the typical documentation tools (e.g. Doxygen or Javadoc)? We could examine issues involving requirements documents [6] for machine learning algorithms.
7. Advice for machine learning people in interfacing with traditional software designers. What are common misunderstandings and things we should be ready to explain?
8. Collaboration on machine learning projects. For instance, platforms that make it easy for engineers to reuse features and code from other teams make every feature engineer much more impactful.
9. Issues with regard to open source in machine learning. Talks involving intellectual property issues in machine learning would also be welcome.
10. Getting data into the software development process is also a possible talk. Handling organization restrictions with regard to security and privacy issues is an important area.
11. Building automatic benchmarking systems. A critical part of machine learning project is to first setup an independent evaluation system to benchmark the current version of the software. This system can ensure that software is not accidentally “peaking” at the test data. Other subtle issues include excessive benchmarking against a test set which could result in overfitting, or not placing any confidence intervals on the benchmarks used. Machine learning competitions can provide some guidance here.
12. Methods for testing and ensuring numerical stability. How do we deal with numerical stability in deployed, or real-time, software systems?
13. Differences between using machine learning in client side vs. server side software. Processing the training set client side, and outside of the designers control, poses many more challenges than only processing the test set on a client machine.
14. Reproducibility: What advice do we have for making machine learning experiments completely reproducible? This is largely an extension of using revision control systems and procedures for logging results.
15. Design patterns: What advice is there for utilizing ideas, for example from Gamma et al. [2], in machine learning projects?

Many of the above items are about utilizing and adapting advice from tradition software development, such as from McConnell [5]. PDF formatted copy of proposal available at: http://goo.gl/1v51aS

WEBSITE: http://www.software4machinelearning.com/