



THE 34TH INTERNATIONAL CONFERENCE ON MACHINE LEARNING

ICML 2017

ICML 2017

SYDNEY, AUSTRALIA
AUGUST 6TH - 11TH

CONFERENCE AT A GLANCE

SUNDAY AUGUST 6TH

Tutorials Session One	8:45 - 11:00 am
Tutorials Session Two	1:00 - 3:15 pm
Tutorials Session Three	3:45 - 6:00 pm
Reception (Level 5, Grand Ballroom)	6:00 - 8:00 pm

MONDAY AUGUST 7TH

Opening Remarks	8:45 am
Invited Talk: Latanya Sweeney (page 12)	9:00 - 10:00 am
Session One	10:30 am - Noon
Session Two	1:30 - 3:00 pm
Session Three	3:30 - 5:00 pm
Session Four	5:15 - 6:45 pm
Poster Session (Level 2, Gallery)	6:45 - 10:00 pm

TUESDAY AUGUST 8TH

Test Of Time Paper	9:00 - 10:00 am
Session One	10:30 am - Noon
Session Two	1:30 - 3:00 pm
Session Three	3:30 - 5:00 pm
Invited Talk: Peter Donnelly (page 12)	5:15 - 6:15 pm
Poster Session (Level 2, Gallery)	6:15 - 10:00 pm

WEDNESDAY AUGUST 9TH

Invited Talk: Raia Hadsell (page 13)	9:00 - 10:00 am
Session One	10:30 am - Noon
Session Two	1:30 - 3:00 pm
Session Three	3:30 - 5:00 pm
Invited Talk: Bernhard Schölkopf (page 13)	5:15 - 6:15 pm
Reception (Level 5, Grand Ballroom)	6:30 - 8:30 pm
Poster Session (Level 2, Gallery)	6:15 - 10:00 pm

THURSDAY AUGUST 10TH

Workshop Sessions	8:30 am - 5:30 pm
	Level 4 & Parkside 1 (Level 2)

FRIDAY AUGUST 11TH

Workshop Sessions	8:30 am - 5:30 pm
	Level 4 & Parkside 1 (Level 2)

Contents

Event Sponsors	2
Conference & Workshop Maps	5
Welcome Letter	6
Sponsor Map & General Information	7
Organizing committee	8
Sunday Tutorial Sessions	9
Invited Speakers	12
Monday Sessions	14
Monday Poster Session	23
Tuesday Sessions	27
Tuesday Poster Session	34
Wednesday Sessions	38
Wed. Poster Session	45
Workshop Schedule	49
Thursday Workshops	50
Friday Workshops	54
Scholar Awards	58
Author Index	59



A Special Thank You To Our Sponsors!

Funding for student travel awards was generously provided by our sponsors, with particular thanks to diamond sponsor NVIDIA. Their exemplary support helped provide travel scholarships so that 217 student researchers could attend ICML to present their research. We are grateful for the support and generosity of our sponsors for helping make ICML a more dynamic and inclusive scientific community.

DIAMOND SPONSOR



NVIDIA is a computer technology company that has pioneered GPU-accelerated computing. The GPU has proven to be unbelievably effective at solving some of the most complex problems in computer science. It started out as an engine for simulating human imagination, conjuring up the amazing virtual worlds of video games and Hollywood films. Today, NVIDIA's GPU simulates human intelligence, running deep learning algorithms and acting as the brain of computers, robots, and self-driving cars that can perceive and understand the world. This is our life's work — to amplify human imagination and intelligence.

PLATINUM SPONSORS



Giving people the power to share and connect requires constant innovation. At Facebook, research permeates everything we do. We believe the most interesting research questions are derived from real world problems. Working on cutting edge research with a practical focus, we push product boundaries every day. At the same time, we publish papers, give talks, and collaborate broadly with the academic community.



Research at Google tackles the most challenging problems in Computer Science and related fields. Being bold and taking risks is essential to what we do, and research teams are embedded throughout Google allowing our discoveries to affect billions of users each day.



JD.com is the largest e-commerce company in China and the largest Chinese retailer, both in terms of revenue. The company strives to offer consumers the best online shopping experience. Through its user-friendly website, native mobile apps, and WeChat and Mobile QQ entry points, JD offers consumers a superior shopping experience. The company has the largest fulfillment infrastructure of any e-commerce company in China. As of March 31, 2017, JD.com operated 7 fulfillment centers and 263 warehouses covering 2,672 counties and districts across China, staffed by its own employees. JD.com is a member of the NASDAQ100 and a Fortune Global 500 company.



At DeepMind, our mission is to solve intelligence and then use that to make the world a better place. Our motivation in all we do is to maximise the positive and transformative impact of AI. We believe that AI should ultimately belong to the world, in order to benefit the many and not the few, and we steadfastly research, publish and implement our work to that end.



When Amazon.com launched in 1995, it was with the mission "to be Earth's most customer-centric company, where customers can find and discover anything they might want to buy online, and endeavors to offer its customers the lowest possible prices." This goal continues today, but Amazon's customers are worldwide now, and have grown to include millions of Consumers, Sellers, Content Creators, and Developers & Enterprises. Each of these groups has different needs, and we always work to meet those needs, innovating new solutions to make things easier, faster, better, and more cost-effective. To learn more about Machine Learning @ Amazon, visit: amazon.jobs/ICML



Intel is democratizing AI innovations by increasing the accessibility of data, tools, training, and intelligent machines, while collaborating across industries to improve society. Intel is committed to driving the future of AI by offering a complete portfolio to deliver end-to-end AI solutions, using Intel technology and programs to help unlock AI's potential, for businesses and society.

GOLD SPONSORS

NETFLIX

Netflix is the world's leading Internet television network with over 93 million members in over 190 countries enjoying more than 125 million hours of TV shows and movies per day. Machine Learning is critical to Netflix and powers important areas such as our personalized recommendation systems, our search algorithms, and how we build our content library.

 Alibaba Group
阿里巴巴集团

ALIBABA GROUP'S MISSION IS TO MAKE IT EASY TO DO BUSINESS ANYWHERE. Our businesses are comprised of core commerce, cloud computing, digital media and entertainment, innovation initiatives and others. Through investee affiliates, we also participate in the logistics and local services sectors. Ant Financial Services Group was originated from Alipay. It is dedicated to creating an open ecosystem, enabling financial institutions to make rapid progress, and providing inclusive financial services to small and micro enterprises and individual consumers.

 蚂蚁金服
ANT FINANCIAL

 Tibra
2006-2016

Celebrating more than a decade in the financial markets sector, Tibra is a quantitative research and investment group that uses its shareholders' capital across a broad range of investment strategies. Founded in Wollongong, NSW, today Tibra spans the globe, employing approximately 100 professionals throughout offices in London, Sydney, Wollongong, Dubai and Hong Kong.

 AIG

American International Group, Inc. (AIG®)'s vision is to become its clients' most valued insurer. For the past 100 years, AIG has been a leading international insurance organisation serving customers in more than 100 countries and jurisdictions. AIG believes that harnessing the power of machine learning and deep learning techniques is essential to go beyond merely generating new insights from data but also to systematically enhance individual human judgement in real business contexts. If you also feel passionate about evidence-based decision making across the world, connect with AIG!

Yandex

Yandex is a technology company that builds intelligent products and services powered by machine learning. Our goal is to help consumers and businesses better navigate the online and offline world. Since 1997, we have delivered world-class, locally relevant search and information services. Additionally, we have developed market-leading on-demand transportation services, navigation products, and other mobile applications for millions of consumers across the globe. Yandex, which has 17 offices worldwide, has been listed on the NASDAQ since 2011.

 SIGOPT

SigOpt is an optimization platform that seamlessly tunes AI and ML model parameters through a state-of-the-art ensemble of Bayesian and global optimization algorithms behind a simple API. This results in captured performance left on the table by conventional techniques and reduces the time and cost for developing and optimizing models. SigOpt can tune any predictive or machine learning model right in place, and the federated API design ensures no proprietary data leaves your premises. Customers can access a web interface or an API to easily integrate SigOpt into any workflow.

 TWO SIGMA

At Two Sigma, we imagine breakthroughs in investment management, insurance and related fields by pushing the boundaries of what open source and proprietary technology can do. In the process, we work to help real people. Our engineers, data scientists and modelers harness data at tremendous scale, using machine learning, distributed computing and other technologies to build powerful predictive models. Come build with us!



Tencent AI Lab

Tencent AI Lab is a leading AI research and application lab of Tencent, China's largest internet company. It was founded in 2016 and backed by 50 world-class research scientists and 200 experienced engineers in China and US. With a vision of "Make AI Everywhere", the Lab is continuously improving AI's capabilities in understanding, decision-making and creativity. Its research focuses on: machine learning, computer vision, speech recognition and natural language processing. To serve the needs of Tencent's core business, its application focuses on: content, game, social and platform AI.

 Maluuba
A Microsoft company

Maluuba, a Microsoft company, is a leader in using deep reinforcement learning to solve language-understanding problems and in training machines to model decision-making capabilities of the human brain. Based in the AI epicenter of Montreal, we operate one of the world's leading research labs dedicated to solving artificial intelligence through natural language understanding. Our team of renowned experts are dedicated to tackling reasoning, decision making and communication in machines through groundbreaking research.

 Microsoft

At Microsoft, we aim to empower every person and every organization on the planet to achieve more. We care deeply about having a global perspective and making a difference in lives and organizations in all corners of the planet. This involves playing a small part in the most fundamental of human activities: Creating tools that enable each of us along our journey to become something more.

 ORACLE

Tackling Today's Biggest Challenges. 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.

 UBER AI Labs

AI that moves the world. Uber's mission is to make transportation as reliable as running water—everywhere, for everyone. At AI Labs, we drive this mission by developing cutting-edge machine learning algorithms that extend the state of the art. By blending a wide variety of approaches across the field, we deliver innovation to application.

 Baidu Research

Baidu Research brings together global research talent to work on AI technologies in areas such as image recognition, speech recognition, high performance computing, natural language processing and deep learning. Baidu Research comprises four labs: the Silicon Valley AI Lab, the Institute of Deep Learning, the Big Data Lab and the Augmented Reality Lab.

SILVER SPONSORS



Recruit Communications Co., Ltd. is dedicated to using technology and marketing communication to develop services that improve people's everyday lives. We are driven by a desire to improve access to information, leveraging both state-of-the-art technology and massive amounts of data in order to create a better world. To fulfill this purpose, we are proud to be actively engaged in research.



Founded in 2007 by leading machine learning scientists, The Voleon Group designs, develops, and implements advanced technology for investment management. We are committed to solving large-scale financial prediction problems with statistical machine learning.



The D. E. Shaw group is a global investment and technology development firm with more than \$41 billion in investment capital as of January 1, 2017, and offices in North America, Europe, and Asia. Since our founding in 1988, our firm has earned an international reputation for successful investing based on innovation, careful risk management, and the quality and depth of our staff.



Cubist Systematic Strategies, the quantitative investing business of Point72 Asset Management, deploys systematic, computer-driven trading strategies across multiple liquid asset classes. The core of our effort is rigorous research into a wide range of market anomalies, fueled by our unparalleled access to a wide range of publicly available data sources.



The Bosch Center for Artificial Intelligence was founded in early 2017. Our goal is to establish cutting-edge AI technology at Bosch. We strive to design and implement AI technology for smart, connected and learning products across all Bosch businesses – with currently 100 experts in Germany, India and the U.S.



Altrics is taking on the ultimate goal of giving everyone access to the power of data-driven health intelligence. Altrics uses state-of-the-art machine learning technologies including deep learning and evidence based expert analysis to provide real-time clinical insights capable of transforming healthcare from reactive to proactive.



Western Digital is an industry-leading provider of storage technologies and solutions that enable people to create, leverage, experience and preserve data. The company addresses ever-changing market needs by providing a full portfolio of compelling, high-quality storage solutions with customer-focused innovation, high efficiency, flexibility and speed. Our products are marketed under the HGST, SanDisk and WD brands to OEMs, distributors, resellers, cloud infrastructure providers and consumers.

BRONZE SPONSORS



Cisco designs and sells products, provides services, and delivers integrated solutions to develop and connect networks around the world. We are focused on helping our customers use IT to enable, differentiate, or fundamentally define their business strategy and drive growth, improve productivity, reduce costs, mitigate risk, and gain a competitive advantage in an increasingly digital world.



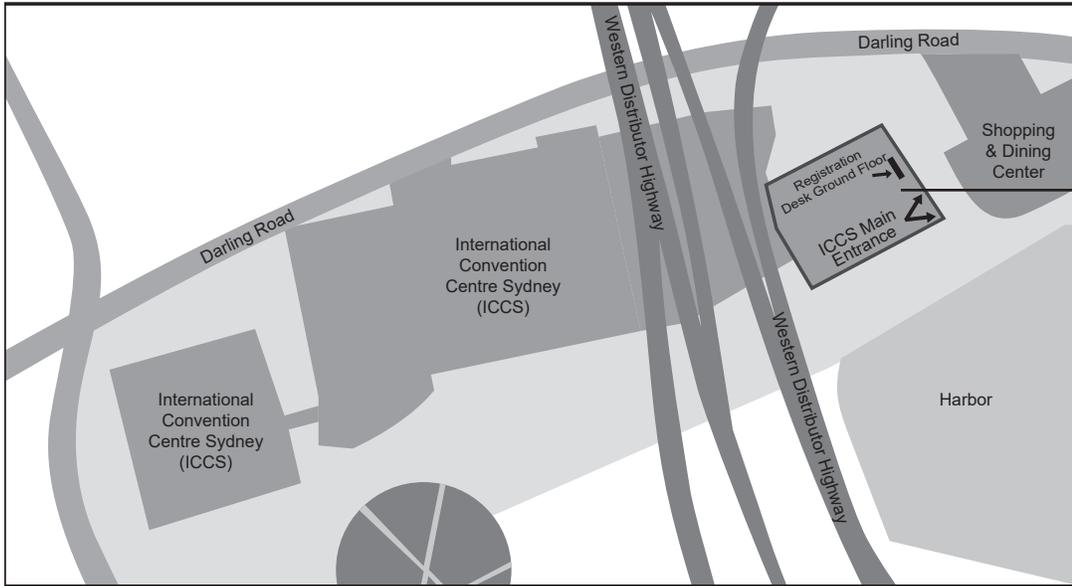
Adobe is the global leader in digital marketing and digital media solutions. Our tools and services allow our customers to create groundbreaking digital content, deploy it across media and devices, measure and optimize it over time and achieve greater business success. We help our customers make, manage, measure and monetize their content across every channel and screen.



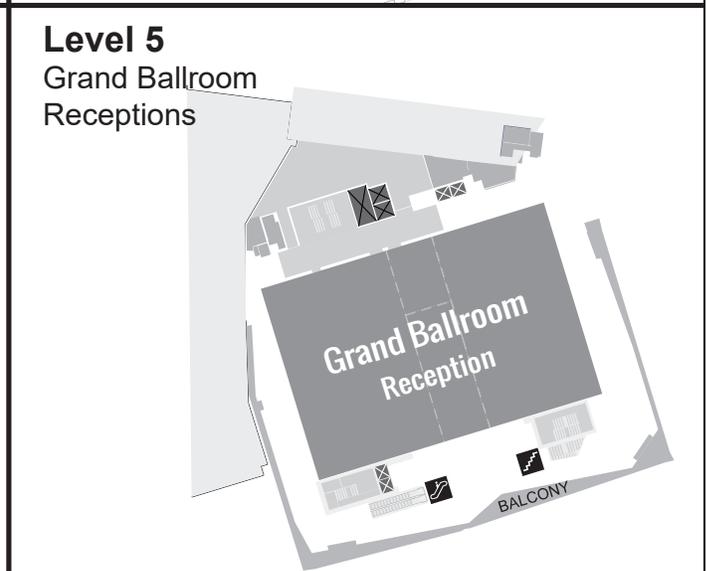
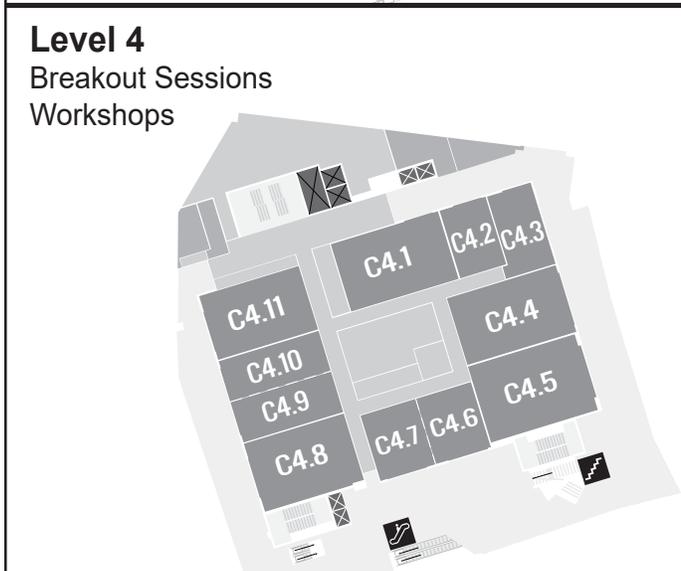
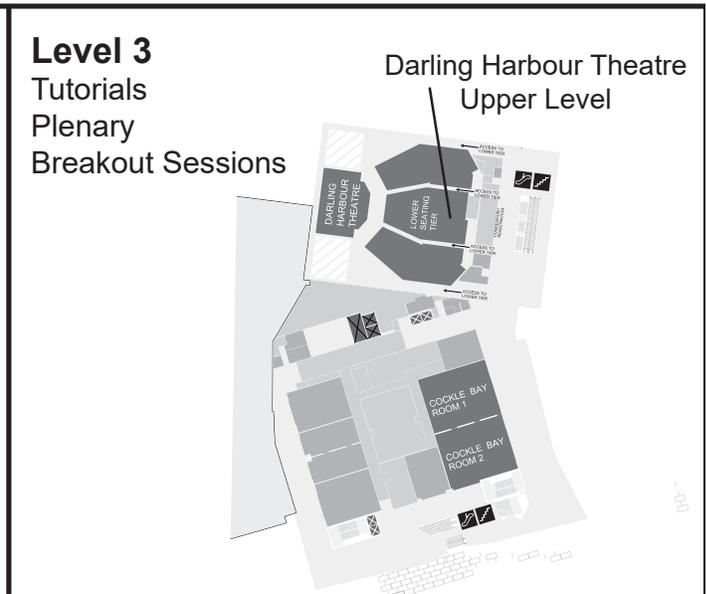
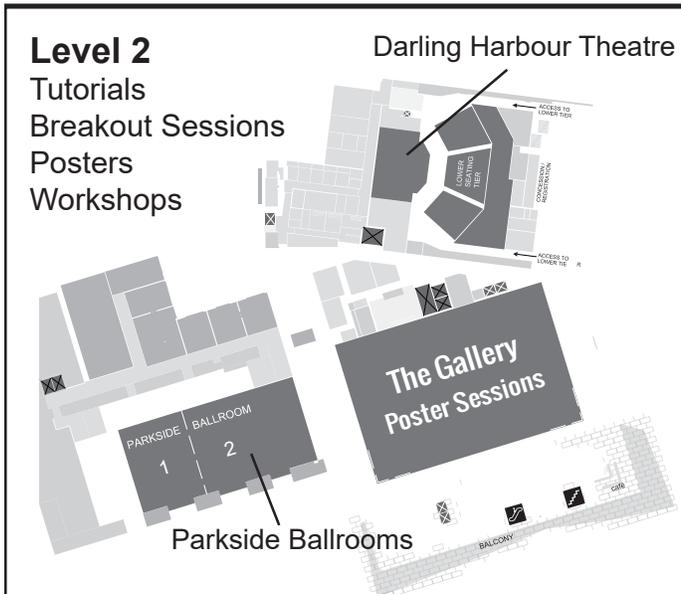
TRI's mission is to use artificial intelligence to improve the quality of human life. We are dedicated to making automobiles safer, more affordable, and more accessible to everyone, regardless of age or ability, and to expanding the benefit of mobility technology beyond automobiles, for example to in-home support of older persons and those with special needs.

ICML 2018:
Stockholm Sweden
Stockholmsmässan,
Tuesday July 10th - Sunday 15th





All Events Will
Take Place
In The Main
Building Here



WELCOME TO ICML 2017 IN SYDNEY

A letter from the Program Chairs

Welcome to Sydney and the 34th International Conference on Machine Learning (ICML 2017)! It is a true pleasure to bring the premiere machine learning conference to Australia and to reunite with you to share the latest breakthroughs in the field. We are excited to hold our meetings on the striking campus of the ICC Sydney which is located in the active precinct of Darling Harbour on Cockle Bay.

Technical Program

The centerpoint of the ICML 2017 conference is the main technical program of contributed papers, talks and posters. This year, ICML continues its rigorous and selective process for identifying impactful and technically sound papers to publish. All in all, ICML 2017 evaluated a record-breaking total of 1676 submissions. Of these, 434 amazing articles (another record!) were accepted for publication and for presentation at the conference. Each accepted contribution is granted both an oral presentation (in one of 9 parallel tracks over three days of meetings) as well a poster presentation (across 3 poster sessions during the evenings). Furthermore, all accepted articles will be published in the Journal of Machine Learning Research (JMLR) under its Workshop and Conference Proceedings series.

The main ICML technical program is interleaved with four world-class keynote talks by invited speakers at the forefront of the field. We are thrilled to host presentations from the following distinguished luminaries: Peter Donnelly from the University of Oxford and Genomics Plc, Latanya Sweeney from Harvard University, Raia Hadsell from DeepMind and Bernhard Schölkopf from the Max Planck Institute for Intelligent Systems.

Tutorials

The main technical program is preceded by a tutorials day featuring 9 tutorial sessions that cover core topics in machine learning today. Attendees will be brought up to speed on the latest advances in subjects such as deep learning, reinforcement learning, distributed learning, interpretable learning, sequence learning, interactive learning, robust algorithms, stochastic optimization, non-convex optimization and application areas (including health care and autonomous vehicles).

Workshops

The main technical program is followed by 22 workshops spread over two days of meetings. These are ideal settings for attendees to see late-breaking work, hear about exploratory topics and pursue collaboration opportunities.

Awards

We will present two best paper awards to honor some of the most promising research from the technical program. We will also present the ICML-2017 Test of Time Award. This award is for the paper from the 2007 ICML conference (held at Oregon State University in Corvallis, USA) that has retrospectively had a significant impact on our field.

Acknowledgements

So many people contributed tremendously to make ICML 2017 a success, we are humbled and deeply thankful for their dedication and their hard work.

Let us take a moment to thank the wonderful service of the 110 distinguished members of the Senior Program Committee and the amazing reviewers they collaborated with. All of them worked hard to give each paper three high-quality reviews as well as an extensive deliberation process that culminated in what was sometimes a very difficult decision. The complete list of all these members of the program committee is available on the ICML website. Through their efforts and world-class expertise, program committee members

help ensure ICML's technical quality and intellectual leadership in the field of machine learning.

We would like to recognize and thank the entire Organizing Committee who put the conference together this year. The Local Chairs, Robert Williamson and Richard Nock were pivotal in securing the location in Sydney and for much of the on-the-ground local work in Australia. The Tutorials Chairs, Alina Beygelzimer and Kamalika Chaudhuri coordinated with many distinguished speakers to put together a fantastic tutorials day to start the conference off. The Workshop Chairs, Fei Sha and Anima Anandkumar curated a wonderful list of 22 workshops. Our Funding Chairs, Ryan Adams and John Cunningham, stunned us with the amazing list of sponsors they were able to enlist and saved the day with the amount of financing they secured. Many thanks to Dan Roy and Iain Murray for their work as Publications Chairs and for putting together the proceedings volume at JMLR. Thanks also to Alan Qi for his assistance in local funding and to Jingrui He for publicizing ICML through various channels. Finally, thanks to our Workflow Chairs, Herke van Hoof and Seth Flaxman, who provided crucially helpful behind-the-scenes work for ICML.

A special thank you is owed to Mary Ellen Perry, the ICML Executive Director, and Lee Campbell who serves as ICML's IT Director. Mary Ellen and Lee have literally been the backbone of this year's conference. Lee maintains ICML's servers, web-pages, registration systems, payment systems, and other workflows. Mary Ellen worked tirelessly to organize logistics, contracts, finances, negotiations, hotel arrangements, travel arrangements, video recording, volunteers, and much more.

We are extremely grateful for the sponsors who helped make ICML a success this year. In particular, their contributions were instrumental in securing registration discounts and travel funds for students who would have otherwise been unable to afford the long journey to Australia. Our thanks go out to Nvidia, Facebook, Google, JD.COM, DeepMind, Amazon, Intel, Netflix, Alibaba, Tibra, AIG, Yandex, SigOpt, Two Sigma, Tencent, Maluuba, Microsoft, Oracle, Uber, Baidu, Recruit Communications, The Voleon Group, D.E. Shaw & Co., Cubist Systematic Strategies, Bosch, Altrics, Western Digital Corporation, Cisco, Adobe and Toyota.

Finally, we want to send a warm thank you to our esteemed IMLS board members and IMLS President Andrew McCallum. Their continued guidance has been crucial this year.

On behalf of all of us at ICML, enjoy the conference!

Doina Precup (Program Co-Chair, ICML 2017)

Yee Whye Teh (Program Co-Chair, ICML 2017)

Tony Jebara (General Chair, ICML 2017)

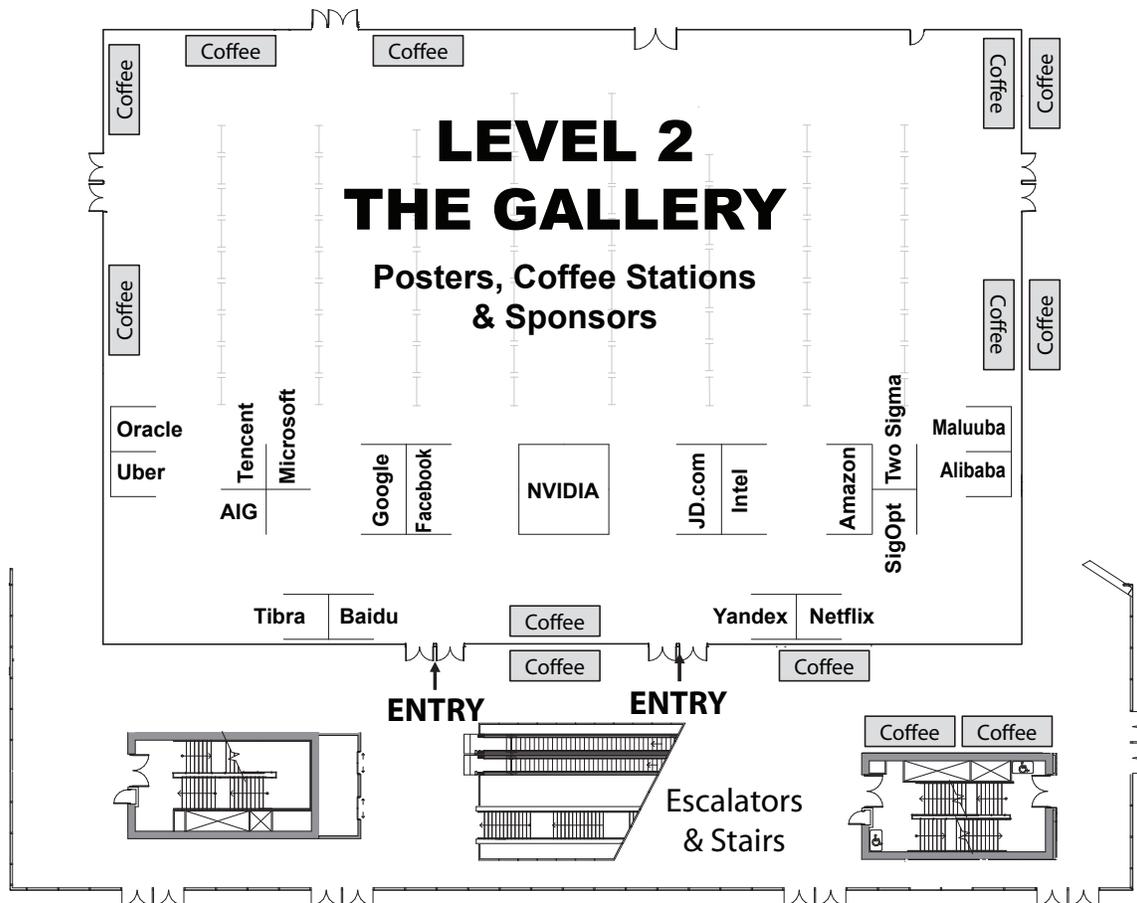
A Warm Welcome From the Local Chairs

Thank you for attending the 34th International Conference on Machine Learning organized by the International Machine Learning Society in Sydney. On behalf of the entire conference organizing committee it is our honor and pleasure to be your hosts. The conference will take place in the vibrant heart of Sydney, which is regularly ranked in the topmost liveable cities in the world.

With some of the most picture-perfect places one can hope for, Sydney has also much to offer outside Darling Harbour. We trust that attendees will find inspiration beyond the conference shore and will find ICML 2017 to be an enjoyable and memorable event.

With best wishes from the Local Chairs,
Richard Nock and Robert C. Williamson

SPONSOR MAP



GENERAL INFORMATION

Attendee Logistics

Event Location

ICML will be held at the International Convention Centre, Sydney Australia.
14 Darling Drive,
Sydney, NSW 2000

Registration

Registration starts at 7 am every day. Registration desk will be on the ground floor near the main entrance on the West side of the building.

Customer Service

Customer Service Desk on the Ground Level of the Convention Centre is open and available to cloak any coats/luggage. Your attendees are more than welcome to cloak their items there until they need to depart for the airport.

Welcome Reception

Sunday, August 6th @ 6 PM
Grand Ballroom - Level 5

Wednesday Reception

Wednesday, August 9th @ 6:30 PM
Grand Ballroom - Level 5

Workshops

All workshops will be held on the 4th floor & Parkside Ballrooms on Thursday and Friday

Poster Sessions

The Gallery - Level 2

Poster Sessions

Monday	6:45 - 10:00 PM
Tuesday	6:15 - 10:00 PM
Wednesday	6:15 - 10:00 PM

Mobile App

- Step 1: Download and install the Whova app from App Store (for iPhones) or Google Play (for Android phones).
- Step 2: Sign up in the app using the email address you registered with.

Now you will be able to:

- View the event agenda and plan your schedule.
- Send in-app messages & exchange contact information (w/profile)
- Receive update notifications
- Access agenda, maps, & directions.

After downloading, sign up on Whova with the email address that you used to RSVP for our event, or sign up using your social media accounts. If you are asked to enter an invitation code to join the event, please use the following invitation code: "icml"

ORGANIZING COMMITTEE



General Chair: Tony Jebara (Columbia & Netflix)

Program Chairs: Doina Precup (McGill University)
Yee Whye Teh (University of Oxford)

Local Chairs: Richard Nock (Data61, The Australian National University and the University of Sydney), Robert Williamson (The Australian National University and Data61)

Tutorial Chairs: Alina Beygelzimer (Yahoo Research), Kamalika Chaudhuri (UCSD)

Workshop Chairs: Anima Anandkumar (UC Irvine)
Fei Sha (USC)

Funding Chairs: John Cunningham (Columbia)
Ryan P Adams (Harvard)

Local Funding Chair: Alan Qi (Ant Financial)

Publication Chairs: Dan Roy (University of Toronto)
Iain Murray (University of Edinburgh)

Workflow Chairs: Herke van Hoof (McGill), Seth Flaxman (Oxford)

Publicity Chair: Jingrui He (Arizona State University)

AREA CHAIRS

Alekh Agarwal (Microsoft Research)
Alex Bouchard-Cote (U. of British Columbia)

Alex Kulesza (Google)
Alexandru Niculescu-Mizil (NEC Laboratories America)
Amr Ahmed (Google)

Andreas Krause (ETH Zurich)
Andrew Dai (Google Brain)
Anna Choromanska (New York U.)
Anthony Lee (University of Warwick)

Arthur Szlam (Facebook)
Ben Marlin (U. Mass Amherst)
Brian Kulis (Boston University)
Charles Sutton (University of Edinburgh)
Cheng Soon Ong (Australian National U.)

Chih-Jen Lin (National Taiwan University)
Corinna Cortes (Google)

Csaba Szepesvari (University of Alberta)
Daniel Hsu (Columbia University)
Danilo Rezende (DeepMind)

Dave Blei (Columbia University)
David Sontag (New York University)
David Duvenaud (University of Toronto)

Dino Sejdinovic (University of Oxford)
Durk Kingma (OpenAI)

Erik Talvitie (Franklin and Marshall College)

Erik Sudderth (Brown University)
Frank Wood (University of Oxford)
Frederik Lindsten (University of Uppsala)
George Dahl (Google Brain)

Guido Sanguinetti (U. of Edinburgh)
Guillaume Obozinski (Ecole des Ponts)
Hal Daume III (University of Maryland)

Honglak Lee (University of Michigan)
Hugo Larochelle (Google)
Ian Goodfellow (OpenAI)

Jacob Abernethy (U. of Michigan)
James Martens (DeepMind)
James Kwok (Hong Kong University of Science and Technology)

Jan Peters (Darmstadt)
Jennifer Dy (Northeastern University)
Joan Bruna

Jon McAuliffe (Voleon Group and University of California at Berkeley)
Jordan Boyd-Graber (University of Maryland)

Joris Mooij (University of Amsterdam)
Jose Miguel Hernandez-Lobato (University of Cambridge)

Julien Mairal (INRIA)
Jun Zhu (Tsinghua University)
Jürgen Schmidhuber (Swiss AI Lab, IDSIA (USI & SUPSI))

Kamalika Chaudhuri (University of California at San Diego)
Karsten Borgwardt (ETH Zurich)

Kate Saenko (Boston University)
Katherine Heller (Duke University)

Kenji Fukumizu (The Institute of Statistical Mathematics)

Kyunghyun Cho (NYU)
Laurens van der Maaten (Facebook AI Research)

Le Song (Georgia Institute of Tech)
Lester Mackey (Microsoft Research)

Lihong Li (Microsoft Research)
Mark Schmidt (University of British Columbia)

Mark Herbster (UCL)
Martin Jaggi (EPFL)
Martin Takac (Lehigh)
Masashi Sugiyama (RIKEN / The University of Tokyo)

Matthew Hoffman (Google Research)
Maya Gupta (Google)

Michalis Titsias (Athens University of Economics and Business)
Mike Osborne (University of Oxford)

Mikhail Belkin (Ohio State University)
Moritz Hardt (Google)

Nando de Freitas (DeepMind)
Nati Srebro (Toyota Technological

Institute at Chicago)
Oriol Vinyals (DeepMind)
Peter Gehler (University of Würzburg)
Peter Bartlett (UC Berkeley)

Phil Long
Philipp Hennig (Max Planck Institute)
Po-Ling Loh (University of Wisconsin)
Pradeep Ravikumar (CMU)

Pushmeet Kohli (Microsoft Research)
Quoc Le (Google Brain)
Raia Hadsell (DeepMind)

Raquel Urtasun (University of Toronto)
Razvan Pascanu (DeepMind)
Ricardo Silva (UCL)

Rich Sutton (University of Alberta)
Robin Evans (University of Oxford)
Roger Grosse (University of Toronto)

Ronan Collobert (Facebook AI Research)
Ruslan Salakhutdinov (CMU)

Ryota Tomioka (Microsoft Research)
Sebastian Bubeck (Microsoft Research)
Sebastian Nowozin (Microsoft Research)

Sergey Levine (Berkeley)
Shakir Mohamed (DeepMind)
Shie Mannor (Technion)

Shimon Whiteson (University of Oxford)
Simon Lacoste-Julien (U. of Montreal)

Sinead Williamson (UT Austin)
Stefanie Jegelka (MIT)
Surya Ganguli (Stanford University)

Thorsten Joachims (Cornell University)
Tim Salimans (OpenAI)
Tong Zhang (Tencent)

Trevor Darrell (UC Berkeley)
Vinayak Rao (Purdue University)
Xuanlong Nguyen (U. of Michigan)

Yann Dauphin (Facebook AI Research)
Yisong Yue (Caltech)
Zaid Harchaoui (U. of Washington)

Zoltan Szabo (École Polytechnique)

Sunday

AUGUST 6TH | TUTORIALS



TUTORIAL SESSION 1 - 8:45 - 11:00 AM

Distributed Deep Learning With MxNet

Alex Smola (Amazon)
Aran Khanna (AWS)

Location: Level 3, Cockle Bay



Software tools for deep learning have progressed significantly recently, making the design of basic computer vision models rather effortless. Distributed inference and sequence models are a different matter. In this tutorial we explain the basic concepts underlying distributed optimization algorithms, such as model and data-parallel strategies, batch sizes and learning rates, and how to configure compute clusters on the fly for deep learning. Moreover, we will illustrate how to build such models for computer vision (e.g. Inception, ResNet) and sequence models (e.g. bidirectional-LSTMs) in MxNet for good scalability.

Machine Learning for Autonomous Vehicles

Raquel Urtasun (Univ. of Toronto)
Andrew Gray (Uber)

Location: Level 2, Parkside 2



The tutorial will cover core machine learning topics for self-driving cars. The objectives are (1) to call to arms of researchers and practitioners to tackle the pressing challenges of autonomous driving; (2) equip participants with enough background to attend the companion

workshop on ML for autonomous vehicles. Machine learning holds the key to solve autonomous driving. Despite recent advances, major problems are far from solved both in terms of fundamental research and engineering challenges.

Interpretable Machine Learning

Been Kim (presenter, Google Brain)
Finale Doshi-Velez (Harvard)

Location: Level 2, Parkside 1



As machine learning systems become ubiquitous, there has been a surge of interest in interpretable machine learning: systems that provide explanation for their outputs. These explanations are often used to qualitatively assess other criteria such as safety or non-discrimination. However, despite the interest in interpretability, there is little consensus on what interpretable machine learning is and how it should be measured. In this talk, we first suggest a definitions of interpretability and describe when interpretability is needed (and when it is not). Then we will review related work, all the way back from classical AI systems to recent efforts for interpretability in deep learning. Finally, we will talk about a taxonomy for rigorous evaluation, and recommendations for researchers. We will end with discussing open questions and concrete problems for new researchers.



TUTORIAL SESSION 2 - 1:00 - 3:15 PM

Recent Advances in Stochastic Convex and Non-Convex Optimization

Zeyuan Allen-Zhu (Microsoft Research)

Location: Level 3, Cockle Bay


In this tutorial, we will provide an accessible and extensive overview on recent advances to optimization methods based on stochastic gradient descent (SGD), for both convex and non-convex tasks. In particular, this tutorial shall try to answer the following questions with theoretical support. How can we properly use momentum to speed up SGD? What is the maximum parallel speedup can we achieve for SGD? When should we use dual or primal-dual approach to replace SGD? What is the difference between coordinate descent (e.g. SDCA) and SGD? How is variance reduction affecting the performance of SGD? Why does the second-order information help us improve the convergence of SGD?

Deep Learning for Health Care Applications: Challenges and Solutions

Yan Liu (USC)

Jimeng Sun (Georgia Tech)

Location: Level 2, Parkside 2


It is widely believed that deep learning and artificial intelligence techniques will fundamentally change health care industries. Even though recent development in deep learning has achieved successes in many applications, such as computer vision, natural language processing, speech recognition and so on, health care applications pose many significantly different challenges to existing deep learning models. Examples include but not are limited to interpretations for prediction, heterogeneity in data, missing value, multi-rate multiresolution data, big and small data, and privacy issues.

In this tutorial, we will discuss a series of problems in health care that can benefit from deep learning models, the challenges as well as recent advances in addressing those. We will also include data sets and demos of working systems.

Deep Reinforcement Learning, Decision Making, and Control

Sergey Levine (UC Berkeley)

Chelsea Finn (UC Berkeley)

Location: Level 2, Parkside 1


Deep learning methods, which combine high-capacity neural network models with simple and scalable training algorithms, have made a tremendous impact across a range of supervised learning domains, including computer vision, speech recognition, and natural language processing. This success has been enabled by the ability of deep networks to capture complex, high-dimensional functions and learn flexible distributed representations. Can this capability be brought to bear on real-world decision making and control problems, where the machine must not only classify complex sensory patterns, but choose actions and reason about their long-term consequences?

Decision making and control problems lack the close supervision present in more classic deep learning applications, and present a number of challenges that necessitate new algorithmic developments. In this tutorial, we will cover the foundational theory of reinforcement and optimal control as it relates to deep reinforcement learning, discuss a number of recent results on extending deep learning into decision making and control, including model-based algorithms, imitation learning, and inverse reinforcement learning, and explore the frontiers and limitations of current deep reinforcement learning algorithms.



TUTORIAL SESSION 3 - 3:30 - 6:00 PM

**Robustness Meets Algorithms
(and Vice-Versa)**

Ankur Moitra (MIT)

Location: Level 2, Parkside 2

In every corner of machine learning and statistics, there is a need for estimators that work not just in an idealized model but even when their assumptions are violated. It turns out that being provably robust and being efficiently computable are often at odds with each other. In even the most basic settings such as robustly computing the mean and covariance, until recently the only known estimators were either hard to compute or could only tolerate a negligible fraction of errors in high-dimensional applications.

In this tutorial, we will survey the exciting recent progress in algorithmic robust statistics. We will give the first provably robust and efficiently computable estimators for several fundamental questions that were thought to be hard, and explain the main insights behind them. We will give practical applications to exploratory data analysis. Finally, we raise some philosophical questions about robustness. It is standard to compare algorithms (especially those with provable guarantees) in terms of their running time and sample complexity. But what frameworks can be used to explore their robustness?

**Sequence-To-Sequence
Modeling with Neural
Networks**

Oriol Vinyals (Google DeepMind)

Navdeep Jaitly (NVIDIA)

Location: Level 2, Parkside 1

Sequence-To-Sequence (Seq2Seq) learning was introduced in 2014, and has since been extensively studied and extended to a large variety of domains. Seq2Seq yields state-of-the-art performance on several applications such as machine translation, image captioning, speech generation, or summarization. In this tutorial, we will survey the basics of this framework, its applications, main algorithmic techniques and future research directions.

**Real World Interactive
Learning**

Alekh Agarwal (Microsoft Research)

John Langford (Microsoft Research)

Location: Level 3, Cockle Bay

This is a tutorial about real-world use of interactive and online learning. We focus on systems for practical applications ranging from recommendation tasks and ad-display, to clinical trials and adaptive decision making in computer systems. There is quite a bit of foundational theory and algorithms from the field of machine learning yet practical use is fraught with several challenges. Success in interactive learning requires a complete learning system which handles exploration, data-flow, logging and real-time updating supporting the core algorithm.

Each potential application also comes with multiple design choices and often do not fit the setting in theory as-is. We cover both foundational principles which have proved practically essential as well as recipes for success from practical experience. After the tutorial, participants should have both a firm understanding of the foundations and the practical ability to deploy and start using such a system in an hour.

INVITED SPEAKERS



Latanya Sweeney
Harvard University

Monday August 7th
9:00 AM
Darling Harbour Theater

How AI Designers will Dictate Our Civic Future

Technology designers are the new policymakers. No one elected them, and most people do not know their names, but the decisions they make when producing the latest gadgets and online innovations dictate the code by which we conduct our daily lives and govern our country. Challenges to the privacy and security of our personal data are part of the first wave of this change; as technology progresses, says Latanya Sweeney, every demographic value and every law comes up for grabs and will likely be redefined by what technology does or does not enable. How will it all fit together or fall apart? Join Sweeney, who after serving as chief technology officer at the U.S. Federal Trade Commission, has been helping others unearth unforeseen consequences and brainstorm on how to engineer the way forward.

Biography:

As Professor of Government and Technology in Residence at Harvard University, my mission is create and use technology to assess and solve societal, political and governance problems, and to teach others how to do the same. On focus area is the scientific study of technology's impact on humankind, and I am the Editor-in-Chief of Technology Science. Another focus area is data privacy, and I am the Director of the Data Privacy Lab at Harvard. There are other foci too.

I was formerly the Chief Technology Officer, also called the Chief Technologist, at the U.S. Federal Trade Commission (FTC). It was a fantastic experience! I thank Chairwoman Ramirez for appointing me. One of my goals was to make it easier for others to work on innovative solutions at the intersection of technology, policy and business. Often, I thought of my past students, who primarily came from computer science or governance backgrounds, and who were highly motivated to change the world. I would like to see society harness their energy and get others thinking about innovative solutions to pressing problems. During my time there, I launched the summer research fellows program and blogged on Tech@FTC to facilitate explorations and ignite brainstorming on FTC-related topics.



Peter Donnelly
Director, Wellcome Trust
Centre, Univ. of Oxford, CEO,
Genomics plc.

Tuesday August 8th
5:15 PM
Darling Harbour Theater

Genomics, big data, and machine learning: understanding the human wiring diagram and driving the healthcare revolution

Each of our cells carries two copies of our genome, the 3bn letters of DNA that serves as their instruction manual. The costs of sequencing (reading) a human genome have decreased by more than six orders of magnitude over the last 10-15 years. Globally, perhaps 100,000 whole genomes have been sequenced, with a clear short-term path to several million. In 10-15 years a billion human genomes will have been sequenced, with many of those sequences linked to extensive information about the individuals, from their medical records and wearable devices. The availability of extensive genetic information linked to information about health outcomes and other traits on very large numbers of individuals presents an extraordinary opportunity. Combining genomic information with biological and health measurements on individuals will improve our ability to assess individual health risks, predict outcomes, and personalise medical treatment. But crucially, and perhaps uniquely, genetics also offers the possibility of unravelling causality amongst otherwise highly correlated features. The resulting much deeper understanding of human biology will have a big impact on drug discovery and healthcare delivery. DNA sequence data from different individuals has a complex correlation structure due to our shared evolutionary history. Inference methods which model these correlations have been very successful to date, but the explosion in the scale and nature of available data will require novel approaches. The talk will illustrate the opportunities and challenges in applying ML and other inference tools to genomic data, by walking through specific examples. No previous knowledge of genetics will be necessary.

Biography:

Donnelly is Director of the Wellcome Trust Centre for Human Genetics and Professor of Statistical Science at the University of Oxford. He grew up in Australia and on graduating from the University of Queensland he studied for a doctorate in Oxford as a Rhodes Scholar. He held professorships at the Universities of London and Chicago before returning to Oxford in 1996. Peter's early research work concerned the development of stochastic models in population genetics, including the coalescent, and then the development of statistical methods for genetic and genomic data. His group developed several widely-used statistical algorithms, including STRUCTURE and PHASE, and, in collaboration with colleagues in Oxford, IMPUTE. His current research focuses on understanding the genetic basis of human diseases, human demographic history, and the mechanisms involved in meiosis and recombination.

Peter played a major role in the HapMap project, and more recently, he chaired the Wellcome Trust Case Control Consortium (WTCCC) and its successor, WTCCC2, a large international collaboration studying the genetic basis of more than 20 common human diseases and conditions in over 60,000 people. He also led WGS500, an Oxford collaboration with Illumina to sequence 500 individuals with a range of clinical conditions to assess the short-term potential for whole genome sequencing in clinical medicine; a precursor to the NHS 100,000 Genomes Project. Peter is a Fellow of the Royal Society and of the Academy of Medical Sciences, and is an Honorary Fellow of the Institute of Actuaries. He has received numerous awards and honours for his research.



Raia Hadsell

DeepMind

**Wednesday August 9th
9:00 AM**

Darling Harbour Theater

Towards Reinforcement Learning in the Real World

Deep reinforcement learning has rapidly grown as a research field with far-reaching potential for artificial intelligence. Large set of ATARI games have been used as the main benchmark domain for many fundamental developments. As the field matures, it is important to develop more sophisticated learning systems with the aim of solving more complex tasks. I will describe some recent research from DeepMind that allows end-to-end learning in challenging environments with real-world variability and complex task structure.

Biography:

Raia Hadsell, a senior research scientist at DeepMind, has worked on deep learning and robotics problems for over 10 years. Her early research developed the notion of manifold learning using Siamese networks, which has been used extensively for invariant feature learning. After completing a PhD with Yann LeCun, which featured a self-supervised deep learning vision system for a mobile robot, her research continued at Carnegie Mellon's Robotics Institute and SRI International, and in early 2014 she joined DeepMind in London to study artificial general intelligence. Her current research focuses on the challenge of continual learning for AI agents and robotic systems. While deep RL algorithms are capable of attaining superhuman performance on single tasks, they cannot transfer that performance to additional tasks, especially if experienced sequentially. She has proposed neural approaches such as policy distillation, progressive nets, and elastic weight consolidation to solve the problem of catastrophic forgetting and improve transfer learning.



Bernhard Schölkopf

Yale University

**Wednesday August 9th
5:15 PM**

Darling Harbour Theater

Causal Learning

In machine learning, we use data to automatically find dependences in the world, with the goal of predicting future observations. Most machine learning methods build on statistics, but one can also try to go beyond this, assaying causal structures underlying statistical dependences. Can such causal knowledge help prediction in machine learning tasks? We argue that this is indeed the case, due to the fact that causal models are more robust to changes that occur in real world datasets. We touch upon the implications of causal models for machine learning tasks such as domain adaptation, transfer learning, and semi-supervised learning.

We also present an application to the removal of systematic errors in astronomical problems such as exoplanet detection.

Biography:

Bernhard Schölkopf's scientific interests are in machine learning and causal inference. He has applied his methods to a number of different application areas, ranging from biomedical problems to computational photography and astronomy. Bernhard has researched at AT&T Bell Labs, at GMD FIRST, Berlin, and at Microsoft Research Cambridge, UK, before becoming a Max Planck director in 2001. He is a member of the German Academy of Sciences (Leopoldina), and has received the J.K. Aggarwal Prize of the International Association for Pattern Recognition, the Max Planck Research Award (shared with S. Thrun), the Academy Prize of the Berlin-Brandenburg Academy of Sciences and Humanities, and the Royal Society Milner Award.

Monday

AUGUST 7TH | SESSIONS



TIME	DESCRIPTION	LOCATION	TIME	DESCRIPTION	LOCATION
8:45 am	Opening remarks	L2, Darling Harbour	3 pm	Break	
9 am	Invited Talk: Latanya Sweeney How AI Designers will Dictate Our Civic Future	L2, Darling Harbour	3:30 pm	SESSION THREE Deep learning 3: Metalearning	L2, Darling Harbour
10 am	Break			Deep generative models 3	L2, Parkside 1
10:30 am	SESSION ONE Deep learning 1: Backdrop	L2, Darling Harbour		Continuous optimization 3	L2, Parkside 2
	Deep generative models 1	L2, Parkside 1		Reinforcement learning 3	L4, C4.5
	Continuous optimization 1	L2, Parkside 2		Probabilistic learning 3	L4, C4.9 & C4.10
	Reinforcement learning 1	L4, C4.5		Online learning 3	L4, C4.1
	Probabilistic learning 1	L4, C4.9 & C4.10		Matrix factorization 2	L4, C4.4
	Online learning 1	L4, C4.1		Deep learning theory 3	L4, C4.8
	Latent feature models	L4, C4.4	5:15 pm	SESSION FOUR Deep learning 4:	
	Deep learning theory 1	L4, C4.8		Learning To Learn	L3, Darling Harbour
	Supervised learning 1	L4, C4.6 & C4.7		Recurrent neural networks 1	L2, Parkside 1
12 pm	LUNCH (On Your Own)			Continuous optimization 4	L2, Parkside 2
1:30 pm	SESSION TWO Deep learning 2:			Reinforcement learning 4	L4, C4.5
	Invariances	L2, Darling Harbour		Bayesian Nonparametrics	L4, C4.9 & C4.10
	Deep generative models 2	L2, Parkside 1		Online learning 4	L4, C4.1
	Continuous optimization 2	L2, Parkside 2		Matrix factorization 3	L4, C4.4
	Reinforcement learning 2	L4, C4.5		Active learning	L4, C4.8
	Probabilistic learning 2	L4, C4.9 & C4.10		Transfer and multitask learning	L4, C4.6 & C4.7
	Online learning 2	L4, C4.1	6:45 pm	Poster Session	L2, Gallery
	Matrix factorization 1	L4, C4.4			
	Deep learning theory 2	L4, C4.8			
	Supervised learning 2	L4, C4.6 & C4.7			



SESSION ONE - 10:30 AM - 12:00 PM

Deep learning 1

Location: Level 2, Darling Harbour

- **Decoupled Neural Interfaces using Synthetic Gradients**
Max Jaderberg, Wojciech Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, Koray Kavukcuoglu
- **Understanding Synthetic Gradients and Decoupled Neural Interfaces**
Wojciech Czarnecki, Grzegorz Świrszcz, Max Jaderberg, Simon Osindero, Oriol Vinyals, Koray Kavukcuoglu
- **meProp: Sparsified Back Propagation for Accelerated Deep Learning with Reduced Overfitting**
Xu SUN, Xuancheng REN, Shuming Ma, Houfeng Wang
- **Learning Important Features Through Propagating Activation Differences**
Avanti Shrikumar, Peyton Greenside, Anshul Kundaje
- **Evaluating the Variance of Likelihood-Ratio Gradient Estimators**
Seiya Tokui, Issei Sato

Deep generative models 1

Location: Level 2, Parkside 1

- **PixelCNN Models with Auxiliary Variables for Natural Image Modeling**
Alexander Kolesnikov, Christoph Lampert
- **Parallel Multiscale Autoregressive Density Estimation**
Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, Nando de Freitas
- **Video Pixel Networks**
Nal Kalchbrenner, Karen Simonyan, Aäron van den Oord, Ivo Danihelka, Oriol Vinyals, Alex Graves, Koray Kavukcuoglu
- **Adversarial Variational Bayes: Unifying Variational Autoencoders and Generative Adversarial Networks**
Lars Mescheder, Sebastian Nowozin, Andreas Geiger
- **Learning Texture Manifolds with the Periodic Spatial GAN**
Urs M Bergmann, Nikolay Jetchev, Roland Vollgraf

Continuous optimization 1

Location: Level 2, Parkside 2

- **Tight Bounds for Approximate Carathéodory and Beyond**
Vahab Mirrokni, Renato Leme, Adrian Vladu, Sam Wong
- **Oracle Complexity of Second-Order Methods for Finite-Sum Problems**
Yossi Arjevani, Ohad Shamir
- **Global optimization of Lipschitz functions**
Cédric Malherbe, Nicolas Vayatis
- **Strong NP-Hardness for Sparse Optimization with Concave Penalty Functions**
Yichen Chen, Dongdong Ge, Mengdi Wang, Zizhuo Wang, Yinyu Ye, Hao Yin
- **Stochastic Convex Optimization: Faster Local Growth Implies Faster Global Convergence**
Yi Xu, Qihang Lin, Tianbao Yang

Reinforcement learning 1

Location: Level 4, C 4.5

- **Robust Adversarial Reinforcement Learning**
Lerrel Pinto, James Davidson, RAHUL Sukthankar, Abhinav Gupta
- **Minimax Regret Bounds for Reinforcement Learning**
Mohammad Gheshlaghi Azar, Ian Osband, Remi Munos
- **Fairness in Reinforcement Learning**
Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Aaron Roth
- **Boosted Fitted Q-Iteration**
Samuele Tosatto, Matteo Pirota, Carlo D'Eramo, Marcello Restelli
- **Why is Posterior Sampling Better than Optimism for Reinforcement Learning?**
Ian Osband, Benjamin Van Roy



SESSION ONE - 10:30 AM - 12:00 PM

Probabilistic learning 1

Location: Level 4, C 4.9 & C 4.10

- **Robust Probabilistic Modeling with Bayesian Data Reweighting**
Yixin Wang, Alp Kucukelbir, David Blei
- **Post-Inference Prior Swapping**
Willie Neiswanger, Eric Xing
- **Evaluating Bayesian Models with Posterior Dispersion Indices**
Alp Kucukelbir, Yixin Wang, David Blei
- **Automatic Discovery of the Statistical Types of Variables in a Dataset**
Isabel Valera, Zoubin Ghahramani
- **Bayesian Models of Data Streams with Hierarchical Power Priors**
Andres Masegosa, Thomas D. Nielsen, Helge Langseth, Dario Ramos-Lopez, Antonio Salmeron, Anders Madsen

Online learning 1

Location: Level 4, C 4.1

- **Multi-objective Bandits: Optimizing the Generalized Gini Index**
Róbert Busa-Fekete, Balazs Szorenyi, Paul Weng, Shie Mannor
- **Online Learning with Local Permutations and Delayed Feedback**
Liran Szlak, Ohad Shamir
- **Model-Independent Online Learning for Influence Maximization**
Sharan Vaswani, Branislav Kveton, Zheng Wen, Mohammad Ghavamzadeh, Laks V.S Lakshmanan, Mark Schmidt
- **Online Learning to Rank in Stochastic Click Models**
Masrour Zoghi, Tomas Tunys, Mohammad Ghavamzadeh, Branislav Kveton, Csaba Szepesvari, Zheng Wen
- **The Sample Complexity of Online One-Class Collaborative Filtering**
Reinhard Heckel, Kannan Ramchandran

Latent feature models

Location: Level 4, C 4.4

- **Communication-efficient Algorithms for Distributed Stochastic Principal Component Analysis**
Dan Garber, Ohad Shamir, Nati Srebro

- **SPLICE: Fully Tractable Hierarchical Extension of ICA with Pooling**
Jun-ichiro Hirayama, Aapo Hyvärinen, Motoaki Kawanabe
- **Latent Feature Lasso**
Ian Yen, Wei-Chen Li, Sung-En Chang, Arun Suggala, Shou-De Lin, Pradeep Ravikumar
- **Online Partial Least Square Optimization: Dropping Convexity for Better Efficiency and Scalability**
Zhehui Chen, Lin Yang, Chris Junchi Li, Tuo Zhao

Deep learning theory 1

Location: Level 4, C 4.8

- **The loss surface of deep and wide neural networks**
Quynh Nguyen, Matthias Hein
- **Neural Taylor Approximations: Convergence and Exploration in Rectifier Networks**
David Balduzzi, Brian McWilliams, Tony Butler-Yeoman
- **Sharp Minima Can Generalize For Deep Nets**
Laurent Dinh, Razvan Pascanu, Samy Bengio, Yoshua Bengio
- **Geometry of Neural Network Loss Surfaces via Random Matrix Theory**
Jeffrey Pennington, Yasaman Bahri
- **The Shattered Gradients Problem: If resnets are the answer, then what is the question?**
David Balduzzi, Marcus Frean, Wan-Duo Ma, Brian McWilliams, Lennox Leary, J.P. Lewis

Supervised learning 1

Location: Level 4, C 4.6 & C 4.7

- **Enumerating Distinct Decision Trees**
Salvatore Ruggieri
- **Simultaneous Learning of Trees and Representations for Extreme Classification and Density Estimation**
Yacine Jernite, Anna Choromanska, David Sontag
- **Resource-efficient Machine Learning in 2 KB RAM for the Internet of Things**
Ashish Kumar, Saurabh Goyal, Manik Varma
- **Multi-Class Optimal Margin Distribution Machine**
Teng Zhang, Zhi-Hua Zhou
- **Kernelized Support Tensor Machines**
Lifang He, Chun-Ta Lu, Guixiang Ma, Shen Wang, Linlin Shen, Philip Yu, Ann Ragin



SESSION TWO - 1:30 PM - 3:00 PM

Deep learning 2: Invariances

Location: Level 2, Darling Harbour Theatre

- **Equivariance Through Parameter-Sharing**
Siamak Ravanbakhsh, Jeff Schneider, Barnabás Póczos
- **Warped Convolutions: Efficient Invariance to Spatial Transformations**
Joao Henriques, Andrea Vedaldi
- **Graph-based Isometry Invariant Representation Learning**
Renata Khasanova, Pascal Frossard
- **Deriving Neural Architectures from Sequence and Graph Kernels**
Tao Lei, Wengong Jin, Regina Barzilay, Tommi Jaakkola
- **Unsupervised Learning by Predicting Noise**
Piotr Bojanowski, Armand Joulin

Deep generative models 2

Location: Level 2, Parkside 1

- **Generalization and Equilibrium in Generative Adversarial Nets (GANs)**
Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, Yi Zhang
- **McGan: Mean and Covariance Feature Matching GAN**
Youssef Mroueh, Tom Sercu, Vaibhava Goel
- **Conditional Image Synthesis with Auxiliary Classifier GANs**
Augustus Odena, Christopher Olah, Jon Shlens
- **Learning to Discover Cross-Domain Relations with Generative Adversarial Networks**
Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jungkwon Lee, Jiwon Kim
- **Wasserstein Generative Adversarial Networks**
Martin Arjovsky, Soumith Chintala, Léon Bottou

Continuous optimization 2

Location: Level 2, Parkside 2

- **GSOS: Gauss-Seidel Operator Splitting Algorithm for Multi-Term Nonsmooth Convex Composite Optimization**
Li Shen, Wei Liu, Ganzhao Yuan, Shiqian Ma
- **Breaking Locality Accelerates Block Gauss-Seidel**
Stephen Tu, Shivaram Venkataraman, Ashia Wilson, Alex Gittens, Michael Jordan, Benjamin Recht

- **Stochastic DCA for the Large-sum of Non-convex Functions Problem and its Application to Group Variable Selection in Classification**
Hoai An Le Thi, Hoai Minh Le, Duy Nhat Phan, Bach Tran
- **Gradient Projection Iterative Sketch for Large-Scale Constrained Least-Squares**
Junqi Tang, Mohammad Golbabaee, Mike E Davies
- **Connected Subgraph Detection with Mirror Descent on SDPs**
Cem Aksoylar, Orecchia Lorenzo, Venkatesh Saligrama

Reinforcement learning 2

Location: Level 4, C 4.5

- **Constrained Policy Optimization**
Joshua Achiam, Dave Held, Aviv Tamar, Pieter Abbeel
- **Reinforcement Learning with Deep Energy-Based Policies**
Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, Sergey Levine
- **Prediction and Control with Temporal Segment Models**
Nikhil Mishra, Pieter Abbeel, Igor Mordatch
- **An Alternative Softmax Operator for Reinforcement Learning**
Kavosh Asadi, Michael L. Littman
- **Fake News Mitigation via Point Process Based Intervention**
Mehrdad Farajtabar, Jiachen Yang, Xiaojing Ye, Huan Xu, Rakshit Trivedi, Elias Khalil, Shuang Li, Le Song, Hongyuan Zha

Probabilistic learning 2

Location: Level 4, C 4.9 & C 4.10

- **Ordinal Graphical Models: A Tale of Two Approaches**
Arun SUGGALA, Eunho Yang, Pradeep Ravikumar
- **Scalable Bayesian Rule Lists**
Hongyu Yang, Cynthia Rudin, Margo Seltzer
- **Learning Determinantal Point Processes with Moments and Cycles**
John C Urschel, Ankur Moitra, Philippe Rigollet, Victor-Emmanuel Brunel
- **Deep Bayesian Active Learning with Image Data**
Yarin Gal, Riashat Islam, Zoubin Ghahramani
- **Bayesian Boolean Matrix Factorisation**
Tammo Rukat, Christopher Holmes, Michalis Titsias, Christopher Yau



SESSION TWO - 1:30 PM - 3:00 PM

Online learning 2

Location: Level 4, C 4.1

- **Efficient Regret Minimization in Non-Convex Games**
Elad Hazan, Karan Singh, Cyril Zhang
- **Identify the Nash Equilibrium in Static Games with Random Payoffs**
Yichi Zhou, Jialian Li, Jun Zhu
- **Follow the Compressed Leader: Faster Online Learning of Eigenvectors and Faster MMWU**
Zeyuan Allen-Zhu, Yuanzhi Li
- **On Kernelized Multi-armed Bandits**
Sayak Ray Chowdhury, Aditya Gopalan
- **Second-Order Kernel Online Convex Optimization with Adaptive Sketching**
Daniele Calandriello, Alessandro Lazaric, Michal Valko

Matrix factorization 1

Location: Level 4, C 4.4

- **Coresets for Vector Summarization with Applications to Network Graphs**
Dan Feldman, Sedat Ozer, Daniela Rus
- **Partitioned Tensor Factorizations for Learning Mixed Membership Models**
Zilong Tan, Sayan Mukherjee
- **On Mixed Memberships and Symmetric Nonnegative Matrix Factorizations**
Xueyu Mao, Purnamrita Sarkar, Deepayan Chakrabarti
- **Nonnegative Matrix Factorization for Time Series Recovery From a Few Temporal Aggregates**
Jiali Mei, Yohann De Castro, Yannig Goude, Georges Hébrail
- **Frame-based Data Factorizations**
Sebastian Mair, AHCÈNE BOUBEKKI, Ulf Brefeld

Deep learning theory 2

Location: Level 4, C 4.8

- **Recovery Guarantees for One-hidden-layer Neural Networks**
Kai Zhong, Zhao Song, Prateek Jain, Peter Bartlett, Inderjit Dhillon
- **Failures of Gradient-Based Deep Learning**
Shaked Shammah, Shai Shalev-Shwartz, Ohad Shamir
- **Analytical Guarantees on Numerical Precision of Deep Neural Networks**
Charbel Sakr, Yongjune Kim, Naresh Shanbhag
- **Follow the Moving Leader in Deep Learning**
Shuai Zheng, James Kwok
- **Theoretical Properties for Neural Networks with Weight Matrices of Low Displacement Rank**
Liang Zhao, Siyu Liao, Yanzhi Wang, Zhe Li, Jian Tang, Bo Yuan

Supervised learning 2

Location: Level 4, C 4.6 & C 4.7

- **Dual Supervised Learning**
Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, Tie-Yan Liu
- **Learning Infinite Layer Networks without the Kernel Trick**
ROI Livni, Daniel Carmon, Amir Globerson
- **Random Fourier Features for Kernel Ridge Regression: Approximation Bounds and Statistical Guarantees**
Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, Amir Zandieh
- **Logarithmic Time One-Against-Some**
Hal Daumé, NIKOS KARAMPATZIAKIS, John Langford, Paul Mineiro
- **Understanding Black-box Predictions via Influence Functions**
Pang Wei Koh, Percy Liang



SESSION THREE - 3:30 PM - 5:00 PM

Deep learning 3: Metalearning

Location: Level 2, Darling Harbour

- **Deep Transfer Learning with Joint Adaptation Networks**
Mingsheng Long, Han Zhu, Jianmin Wang, Michael Jordan
- **Meta Networks**
Tsendsuren Munkhdalai, Hong Yu
- **SplitNet: Learning to Semantically Split Deep Networks for Parameter Reduction and Model Parallelization**
Juyong Kim, Yookoon Park, Gunhee Kim, Sung Ju Hwang
- **Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks**
Chelsea Finn, Pieter Abbeel, Sergey Levine
- **AdaNet: Adaptive Structural Learning of Artificial Neural Networks**
Corinna Cortes, Xavi Gonzalvo, Vitaly Kuznetsov, Mehryar Mohri, Scott Yang

Deep generative models 3

Location: Level 2, Parkside 1

- **Learning Hierarchical Features from Deep Generative Models**
Shengjia Zhao, Jiaming Song, Stefano Ermon
- **Bottleneck Conditional Density Estimation**
Rui Shu, Hung Bui, Mohammad Ghavamzadeh
- **Learning Deep Latent Gaussian Models with Markov Chain Monte Carlo**
Matthew Hoffman
- **Zero-Inflated Exponential Family Embeddings**
Liping Liu, David Blei

Continuous optimization 3

Location: Level 2, Parkside 2

- **No First Presentation**
- **Exploiting Strong Convexity from Data with Primal-Dual First-Order Algorithms**
Jialei Wang, Lin Xiao
- **Doubly Greedy Primal-Dual Coordinate Descent for Sparse Empirical Risk Minimization**
Qi Lei, Ian Yen, Chao-Yuan Wu, Inderjit Dhillon, Pradeep Ravikumar

- **A Richer Theory of Convex Constrained Optimization with Reduced Projections and Improved Rates**
Tianbao Yang, Qihang Lin, Lijun Zhang
- **Convex Phase Retrieval without Lifting via PhaseMax**
Tom Goldstein, Christoph Studer

Reinforcement learning 3

Location: Level 4, C 4.5

- **Curiosity-driven Exploration by Self-supervised Prediction**
Deepak Pathak, Pulkit Agrawal, Alexei Efros, Prof. Darrell
- **Interactive Learning from Policy-Dependent Human Feedback**
James MacGlashan, Mark Ho, Robert Loftin, Bei Peng, Guan Wang, David L Roberts, Matthew E. Taylor, Michael L. Littman
- **End-to-End Differentiable Adversarial Imitation Learning**
Nir Baram, Oron Anshel, Itai Caspi, Shie Mannor
- **Learning in POMDPs with Monte Carlo Tree Search**
Sammie Katt, Frans A Oliehoek, Chris Amato
- **DARLA: Improving Zero-Shot Transfer in Reinforcement Learning**
Irina Higgins, Arka Pal, Andrei A Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, Alexander Lerchner

Probabilistic learning 3

Location: Level 4, C 4.9 & C 4.10

- **Learning the Structure of Generative Models without Labeled Data**
Stephen Bach, Bryan He, Alexander J Ratner, Christopher Re
- **Learning to Discover Sparse Graphical Models**
Eugene Belilovsky, Kyle Kastner, Gael Varoquaux, Matthew B Blaschko
- **Local-to-Global Bayesian Network Structure Learning**
Tian Gao, Kshitij Fadnis, Murray Campbell
- **Composing Tree Graphical Models with Persistent Homology Features for Clustering Mixed-Type Data**
XIUYAN NI, Novi Quadrianto, Yusu Wang, Chao Chen
- **On Relaxing Determinism in Arithmetic Circuits**
Arthur Choi, Adnan Darwiche



SESSION THREE - 3:30 PM - 5:00 PM

Online learning 3

Location: Level 4, C 4.1

- **Dueling Bandits with Weak Regret**
Bangrui Chen, Peter Frazier
- **On Context-Dependent Clustering of Bandits**
Claudio Gentile, Shuai Li, Puru Kar, Alexandros Karatzoglou, Giovanni Zappella, Evans Etrua Howard
- **Provably Optimal Algorithms for Generalized Linear Contextual Bandits**
Lihong Li, Yu Lu, Dengyong Zhou
- **Safety-Aware Algorithms for Adversarial Contextual Bandit**
Wen Sun, Debadeepta Dey, Ashish Kapoor
- **Adaptive Multiple-Arm Identification**
Jiecao (Jack) Chen, Xi Chen, Qin Zhang, Yuan Zhou

Matrix factorization 2

Location: Level 4, C 4.4

- **Nearly Optimal Robust Matrix Completion**
Yeshwanth Cherapanamjeri, Prateek Jain, Kartik Gupta
- **Provable Alternating Gradient Descent for Non-negative Matrix Factorization with Strong Correlations**
Yuanzhi Li, Yingyu Liang
- **No Spurious Local Minima in Nonconvex Low Rank Problems: A Unified Geometric Analysis**
Rong Ge, Chi Jin, Yi Zheng
- **Coherence Pursuit: Fast, Simple, and Robust Subspace Recovery**
Mostafa Rahmani, George Atia
- **Tensor Decomposition with Smoothness**
Masaaki Imaizumi, Kohei Hayashi

Deep learning theory 3

Location: Level 4, **C4.8**

- **Globally Optimal Gradient Descent for a ConvNet with Gaussian Inputs**
Alon Brutzkus, Amir Globerson
- **Convexified Convolutional Neural Networks**
Yuchen Zhang, Percy Liang, Martin Wainwright
- **On the Expressive Power of Deep Neural Networks**
Maithra Raghu, Ben Poole, Surya Ganguli, Jon Kleinberg, Jascha Sohl-Dickstein
- **Depth-Width Tradeoffs in Approximating Natural Functions With Neural Networks**
Itay Safran, Ohad Shamir

Semisupervised and curriculum learning

Location: Level 4, C 4.6 & C 4.7

- **Re-revisiting Learning on Hypergraphs: Confidence Interval and Subgradient Method**
Chenzi Zhang, Shuguang Hu, Zhihao Gavin Tang, Hubert Chan
- **Self-Paced Co-training**
Fan Ma, Deyu Meng, Qi Xie, Zina Li, Xuanyi Dong
- **Semi-Supervised Classification Based on Classification from Positive and Unlabeled Data**
Tomoya Sakai, Marthinus C du Plessis, Gang Niu, Masashi Sugiyama
- **Iterative Machine Teaching**
Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda Smith, Jim Rehg, Le Song
- **Automated Curriculum Learning for Neural Networks**
Alex Graves, Marc Bellemare, Jacob Menick, Remi Munos, Koray Kavukcuoglu



SESSION FOUR - 5:15 PM - 6:45 PM

Deep learning 4: learning to learn

Location: Level 2, Darling Harbour

- **Learning to Learn without Gradient Descent by Gradient Descent**
Yutian Chen, Matthew Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Tim Lillicrap, Matthew Botvinick, Nando de Freitas
- **Learned Optimizers that Scale and Generalize**
Olga Wichrowska, Niru Maheswaranathan, Matthew Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Nando de Freitas, Jascha Sohl-Dickstein
- **Learning Gradient Descent: Better Generalization and Longer Horizons**
Kaifeng Lv, Shunhua Jiang, Jian Li
- **Learning Algorithms for Active Learning**
Philip Bachman, Alessandro Sordoni, Adam Trischler

Recurrent neural networks 1

Location: Level 2, Parkside 1

- **Attentive Recurrent Comparators**
Pranav Shyam, Shubham Gupta, Ambedkar Dukkipati
- **State-Frequency Memory Recurrent Neural Networks**
Hao Hu, Guo-Jun Qi
- **Delta Networks for Optimized Recurrent Network Computation**
Daniel Neil, Jun Lee, Tobi Delbruck, Shih-Chii Liu
- **Recurrent Highway Networks**
Julian Zilly, Rupesh Srivastava, Jan Koutnik, Jürgen Schmidhuber

Continuous optimization 4

Location: Level 2, Parkside 2

- **A Semismooth Newton Method for Fast, Generic Convex Programming**
Alnur Ali, Eric Wong, Zico Kolter
- **Approximate Newton Methods and Their Local Convergence**
Haishan Ye, Luo Luo, Zhihua Zhang
- **Stochastic Adaptive Quasi-Newton Methods for Minimizing Expected Values**
Chaoxu Zhou, Wenbo Gao, Donald Goldfarb

- **Practical Gauss-Newton Optimisation for Deep Learning**

Alex Botev, Hippolyt Ritter, David Barber

- **Tensor Balancing on Statistical Manifold**

Mahito Sugiyama, Hiroyuki Nakahara, Koji Tsuda

Reinforcement learning 4

Location: Level 4, C 4.5

- **Unifying task specification in reinforcement learning**
Martha White
- **A Distributional Perspective on Reinforcement Learning**
Marc Bellemare, Will Dabney, Remi Munos
- **Hierarchy Through Composition with Multitask LMDPs**
Andrew Saxe, Adam Earle, Benjamin Rosman
- **A Laplacian Framework for Option Discovery in Reinforcement Learning**
Marlos C. Machado, Marc Bellemare, Michael Bowling
- **Modular Multitask Reinforcement Learning with Policy Sketches**
Jacob Andreas, Dan Klein, Sergey Levine

Bayesian Nonparametrics

Location: Level 4, C 4.9 & C 4.10

- **An Infinite Hidden Markov Model With Similarity-Biased Transitions**
Colin Dawson, Bill Huang, Clayton T. Morrison
- **Batched High-dimensional Bayesian Optimization via Structural Kernel Learning**
Zi Wang, Chengtao Li, Stefanie Jegelka, Pushmeet Kohli
- **From Patches to Images: A Nonparametric Generative Model**
Geng Ji, Michael C. Hughes, Erik Sudderth
- **Fast Bayesian Intensity Estimation for the Permenental Process**
Christian Walder, Adrian N Bishop
- **A Birth-Death Process for Feature Allocation**
Konstantina Palla, David Knowles, Zoubin Ghahramani



SESSION FOUR - 5:15 PM - 6:45 PM

Online learning 4

Location: Level 4, C 4.1

- **Efficient Online Bandit Multiclass Learning with $O\sqrt{T}$ Regret**
Alina Beygelzimer, Francesco Orabona, Chicheng Zhang
- **Active Learning for Accurate Estimation of Linear Models**
Carlos Riquelme Ruiz, Mohammad Ghavamzadeh, Alessandro Lazaric
- **Adaptive Feature Selection: Computationally Efficient Online Sparse Linear Regression under RIP**
Satyen Kale, Zohar Karnin, Tengyuan Liang, David Pal
- **Emulating the Expert: Inverse Optimization through Online Learning**
Sebastian Pokutta, Andreas Bärman, Oskar Schneider
- **Variants of RMSProp and Adagrad with Logarithmic Regret Bounds**
Mahesh Chandra Mukkamala, Matthias Hein

Matrix factorization 3

Location: Level 4, C 4.4

- **Orthogonalized ALS: A Theoretically Principled Tensor Decomposition Algorithm for Practical Use**
Vatsal Sharan, Gregory Valiant
- **Tensor Decomposition via Simultaneous Power Iteration**
Poan Wang, Chi-Jen Lu
- **A Unified Variance Reduction-Based Framework for Nonconvex Low-Rank Matrix Recovery**
Lingxiao Wang, Xiao Zhang, Quanquan Gu
- **An Efficient, Sparsity-Preserving, Online Algorithm for Low-Rank Approximation**
Dave Anderson, Ming Gu
- **Algorithms for ϵ p Low-Rank Approximation**
Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, David Woodruff

Active learning

Location: Level 4, C 4.8

- **Efficient Nonmyopic Active Search**
Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, Roman Garnett
- **Leveraging Union of Subspace Structure to Improve Constrained Clustering**
John Lipor, Laura Balzano
- **Active Heteroscedastic Regression**
Kamalika Chaudhuri, Prateek Jain, Nagarajan Natarajan
- **Active Learning for Cost-Sensitive Classification**
Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, Hal Daumé III, John Langford
- **Diameter-Based Active Learning**
Chris Tosh, Sanjoy Dasgupta

Transfer and multitask learning

Location: Level 4, C 4.6 & C 4.7

- **Asymmetric Tri-training for Unsupervised Domain Adaptation**
Saito Kuniaki, Yoshitaka Ushiku, Tatsuya Harada
- **Source-Target Similarity Modelings for Multi-Source Transfer Gaussian Process Regression**
PENGFEI WEI, Ramon Sagarna, Yiping Ke, yEW ONG, CHI GOH
- **Multi-task Learning with Labeled and Unlabeled Tasks**
Anastasia Pentina, Christoph Lampert
- **Schema Networks: Zero-shot Transfer with a Generative Causal Model of Intuitive Physics**
Ken Kansky, Tom Silver, David A Mély, Mo Eldawy, Miguel Lazaro-Gredilla, Xinghua Lou, Nimrod Dorfman, Szymon Sidor, Scott Phoenix, Dileep George
- **Risk Bounds for Transferring Representations With and Without Fine-Tuning**
Daniel McNamara, Nina Balcan



- #1 Decoupled Neural Interfaces using Synthetic Gradients**
Max Jaderberg, Wojciech Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, Koray Kavukcuoglu
- #2 PixelCNN Models with Auxiliary Variables for Natural Image Modeling**
Alexander Kolesnikov, Christoph Lampert
- #3 Tight Bounds for Approximate Carathéodory and Beyond**
Vahab Mirrokni, Renato Leme, Adrian Vladu, Sam Wong
- #4 Robust Adversarial Reinforcement Learning**
Lerrel Pinto, James Davidson, RAHUL Sukthankar, Abhinav Gupta
- #5 Robust Probabilistic Modeling with Bayesian Data Reweighting**
Yixin Wang, Alp Kucukelbir, David Blei
- #6 Multi-objective Bandits: Optimizing the Generalized Gini Index**
Róbert Busa-Fekete, Balazs Szorenyi, Paul Weng, Shie Mannor
- #7 Communication-efficient Algorithms for Distributed Stochastic Principal Component Analysis**
Dan Garber, Ohad Shamir, Nati Srebro
- #8 Enumerating Distinct Decision Trees**
Salvatore Ruggieri
- #9 Understanding Synthetic Gradients and Decoupled Neural Interfaces**
Wojciech Czarnecki, Grzegorz Świrszcz, Max Jaderberg, Simon Osindero, Oriol Vinyals, Koray Kavukcuoglu
- #10 Parallel Multiscale Autoregressive Density Estimation**
Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, Nando de Freitas
- #11 Oracle Complexity of Second-Order Methods for Finite-Sum Problems**
Yossi Arjevani, Ohad Shamir
- #12 Minimax Regret Bounds for Reinforcement Learning**
Mohammad Gheshlaghi Azar, Ian Osband, Remi Munos
- #13 Post-Inference Prior Swapping**
Willie Neiswanger, Eric Xing
- #14 Online Learning with Local Permutations and Delayed Feedback**
Liran Szlak, Ohad Shamir
- #15 SPLICE: Fully Tractable Hierarchical Extension of ICA with Pooling**
Jun-ichiro Hirayama, Aapo Hyvärinen, Motoaki Kawanabe
- #16 Simultaneous Learning of Trees and Representations for Extreme Classification and Density Estimation**
Yacine Jernite, Anna Choromanska, David Sontag
- #17 meProp: Sparsified Back Propagation for Accelerated Deep Learning with Reduced Overfitting**
Xu SUN, Xuancheng REN, Shuming Ma, Houfeng Wang
- #18 Video Pixel Networks**
Nal Kalchbrenner, Karen Simonyan, Aäron van den Oord, Ivo Danihelka, Oriol Vinyals, Alex Graves, Koray Kavukcuoglu
- #19 Global optimization of Lipschitz functions**
Cédric Malherbe, Nicolas Vayatis
- #20 Fairness in Reinforcement Learning**
Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Aaron Roth
- #21 Evaluating Bayesian Models with Posterior Dispersion Indices**
Alp Kucukelbir, Yixin Wang, David Blei
- #22 Model-Independent Online Learning for Influence Maximization**
Sharan Vaswani, Branislav Kveton, Zheng Wen, Mohammad Ghavamzadeh, Laks V.S Lakshmanan, Mark Schmidt
- #23 Latent Feature Lasso**
Ian Yen, Wei-Chen Li, Sung-En Chang, Arun Suggala, Shou-De Lin, Pradeep Ravikumar
- #24 Resource-efficient Machine Learning in 2 KB RAM for the Internet of Things**
Ashish Kumar, Saurabh Goyal, Manik Varma
- #25 Learning Important Features Through Propagating Activation Differences**
Avanti Shrikumar, Peyton Greenside, Anshul Kundaje
- #26 Adversarial Variational Bayes: Unifying Variational Autoencoders and Generative Adversarial Networks**
Lars Mescheder, Sebastian Nowozin, Andreas Geiger
- #27 Strong NP-Hardness for Sparse Optimization with Concave Penalty Functions**
Yichen Chen, Dongdong Ge, Mengdi Wang, Zizhuo Wang, Yinyu Ye, Hao Yin
- #28 Boosted Fitted Q-Iteration**
Samuele Tosatto, Matteo Pirotta, Carlo D'Eramo, Marcello Restelli
- #29 Automatic Discovery of the Statistical Types of Variables in a Dataset**
Isabel Valera, Zoubin Ghahramani
- #30 Online Learning to Rank in Stochastic Click Models**
Masrour Zoghi, Tomas Tunys, Mohammad Ghavamzadeh, Branislav Kveton, Csaba Szepesvari, Zheng Wen
- #31 Online Partial Least Square Optimization: Dropping Convexity for Better Efficiency and Scalability**
Zhehui Chen, Lin Yang, Chris Junchi Li, Tuo Zhao
- #32 Multi-Class Optimal Margin Distribution Machine**
Teng Zhang, Zhi-Hua Zhou
- #33 Evaluating the Variance of Likelihood-Ratio Gradient Estimators**
Seiya Tokui, Issei Sato
- #34 Learning Texture Manifolds with the Periodic Spatial GAN**
Urs M Bergmann, Nikolay Jetchev, Roland Vollgraf



- #35 Stochastic Convex Optimization: Faster Local Growth Implies Faster Global Convergence**
Yi Xu, Qihang Lin, Tianbao Yang
- #36 Why is Posterior Sampling Better than Optimism for Reinforcement Learning?**
Ian Osband, Benjamin Van Roy
- #37 Bayesian Models of Data Streams with Hierarchical Power Priors**
Andres Masegosa, Thomas D. Nielsen, Helge Langseth, Dario Ramos-Lopez, Antonio Salmeron, Anders Madsen
- #38 The Sample Complexity of Online One-Class Collaborative Filtering**
Reinhard Heckel, Kannan Ramchandran
- #39 Kernelized Support Tensor Machines**
Lifang He, Chun-Ta Lu, Guixiang Ma, Shen Wang, Linlin Shen, Philip Yu, Ann Ragin
- #40 Equivariance Through Parameter-Sharing**
Siamak Ravanbakhsh, Jeff Schneider, Barnabás Póczos
- #41 Generalization and Equilibrium in Generative Adversarial Nets (GANs)**
Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, Yi Zhang
- #42 GSOS: Gauss-Seidel Operator Splitting Algorithm for Multi-Term Nonsmooth Convex Composite Optimization**
Li Shen, Wei Liu, Ganzhao Yuan, Shiqian Ma
- #43 Constrained Policy Optimization**
Joshua Achiam, Dave Held, Aviv Tamar, Pieter Abbeel
- #44 Ordinal Graphical Models: A Tale of Two Approaches**
Arun SUGGALA, Eunho Yang, Pradeep Ravikumar
- #45 Efficient Regret Minimization in Non-Convex Games**
Elad Hazan, Karan Singh, Cyril Zhang
- #46 Coresets for Vector Summarization with Applications to Network Graphs**
Dan Feldman, Sedat Ozer, Daniela Rus
- #47 Recovery Guarantees for One-hidden-layer Neural Networks**
Kai Zhong, Zhao Song, Prateek Jain, Peter Bartlett, Inderjit Dhillon
- #48 Dual Supervised Learning**
Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, Tie-Yan Liu
- #49 Warped Convolutions: Efficient Invariance to Spatial Transformations**
Joao Henriques, Andrea Vedaldi
- #50 McGAN: Mean and Covariance Feature Matching GAN**
Youssef Mroueh, Tom Sercu, Vaibhava Goel
- #51 Breaking Locality Accelerates Block Gauss-Seidel**
Stephen Tu, Shivaram Venkataraman, Ashia Wilson, Alex Gittens, Michael Jordan, Benjamin Recht
- #52 Reinforcement Learning with Deep Energy-Based Policies**
Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, Sergey Levine
- #53 Scalable Bayesian Rule Lists**
Hongyu Yang, Cynthia Rudin, Margo Seltzer
- #54 Identify the Nash Equilibrium in Static Games with Random Payoffs**
Yichi Zhou, Jialian Li, Jun Zhu
- #55 Partitioned Tensor Factorizations for Learning Mixed Membership Models**
Zilong Tan, Sayan Mukherjee
- #56 Failures of Gradient-Based Deep Learning**
Shaked Shammah, Shai Shalev-Shwartz, Ohad Shamir
- #57 Learning Infinite Layer Networks without the Kernel Trick**
ROI Livni, Daniel Carmon, Amir Globerson
- #58 Graph-based Isometry Invariant Representation Learning**
Renata Khasanova, Pascal Frossard
- #59 Conditional Image Synthesis with Auxiliary Classifier GANs**
Augustus Odena, Christopher Olah, Jon Shlens
- #60 Stochastic DCA for the Large-sum of Non-convex Functions Problem and its Application to Group Variable Selection in Classification**
Hoai An Le Thi, Hoai Minh Le, Duy Nhat Phan, Bach Tran
- #61 Prediction and Control with Temporal Segment Models**
Nikhil Mishra, Pieter Abbeel, Igor Mordatch
- #62 Learning Determinantal Point Processes with Moments and Cycles**
John C Urschel, Ankur Moitra, Philippe Rigollet, Victor-Emmanuel Brunel
- #63 Follow the Compressed Leader: Faster Online Learning of Eigenvectors and Faster MMWU**
Zeyuan Allen-Zhu, Yuanzhi Li
- #64 On Mixed Memberships and Symmetric Nonnegative Matrix Factorizations**
Xueyu Mao, Purnamrita Sarkar, Deepayan Chakrabarti
- #65 Analytical Guarantees on Numerical Precision of Deep Neural Networks**
Charbel Sakr, Yongjune Kim, Naresh Shanbhag
- #66 Random Fourier Features for Kernel Ridge Regression: Approximation Bounds and Statistical Guarantees**
Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, Amir Zandieh
- #67 Deriving Neural Architectures from Sequence and Graph Kernels**
Tao Lei, Wengong Jin, Regina Barzilay, Tommi Jaakkola
- #68 Learning to Discover Cross-Domain Relations with Generative Adversarial Networks**
Taeksoo Kim, Moon-su Cha, Hyunsoo Kim, Jungkwon Lee, Jiwon Kim
- #69 Gradient Projection Iterative Sketch for Large-Scale Constrained Least-Squares**
Junqi Tang, Mohammad Golbabaee, Mike E Davies
- #70 An Alternative Softmax Operator for Reinforcement Learning**
Kavosh Asadi, Michael L. Littman
- #71 Deep Bayesian Active Learning with Image Data**
Yarin Gal, Riashat Islam, Zoubin Ghahramani



- #72 On Kernelized Multi-armed Bandits**
Sayak Ray Chowdhury, Aditya Gopalan
- #73 Nonnegative Matrix Factorization for Time Series Recovery From a Few Temporal Aggregates**
Jiali Mei, Yohann De Castro, Yannig Goude, Georges Hébrail
- #74 Follow the Moving Leader in Deep Learning**
Shuai Zheng, James Kwok
- #75 Logarithmic Time One-Against-Some**
Hal Daumé, NIKOS KARAMPATZIAKIS, John Langford, Paul Mineiro
- #76 Unsupervised Learning by Predicting Noise**
Piotr Bojanowski, Armand Joulin
- #77 Wasserstein Generative Adversarial Networks**
Martin Arjovsky, Soumith Chintala, Léon Bottou
- #78 Connected Subgraph Detection with Mirror Descent on SDPs**
Cem Aksoylar, Orecchia Lorenzo, Venkatesh Saligrama
- #79 Fake News Mitigation via Point Process Based Intervention**
Mehrdad Farajtabar, Jiachen Yang, Xiaojing Ye, Huan Xu, Rakshit Trivedi, Elias Khalil, Shuang Li, Le Song, Hongyuan Zha
- #80 Bayesian Boolean Matrix Factorisation**
Tammo Rukat, Christopher Holmes, Michalis Titsias, Christopher Yau
- #81 Second-Order Kernel Online Convex Optimization with Adaptive Sketching**
Daniele Calandriello, Alessandro Lazaric, Michal Valko
- #82 Frame-based Data Factorizations**
Sebastian Mair, AHCÈNE BOUBEKKE, Ulf Brefeld
- #83 Theoretical Properties for Neural Networks with Weight Matrices of Low Displacement Rank**
Liang Zhao, Siyu Liao, Yanzhi Wang, Zhe Li, Jian Tang, Bo Yuan
- #84 Understanding Black-box Predictions via Influence Functions**
Pang Wei Koh, Percy Liang
- #85 Deep Transfer Learning with Joint Adaptation Networks**
Mingsheng Long, Han Zhu, Jianmin Wang, Michael Jordan
- #86 Learning Hierarchical Features from Deep Generative Models**
Shengjia Zhao, Jiaming Song, Stefano Ermon
- #87 Prox-PDA: The Proximal Primal-Dual Algorithm for Fast Distributed Nonconvex Optimization and Learning Over Networks**
Mingyi Hong, Davood Hajinezhad, Ming-Min Zhao
- #88 Curiosity-driven Exploration by Self-supervised Prediction**
Deepak Pathak, Pulkit Agrawal, Alexei Efros, Prof. Darrell
- #89 Learning the Structure of Generative Models without Labeled Data**
Stephen Bach, Bryan He, Alexander J Ratner, Christopher Re
- #90 Dueling Bandits with Weak Regret**
Bangrui Chen, Peter Frazier
- #91 Nearly Optimal Robust Matrix Completion**
Yeshwanth Cherapanamjeri, Prateek Jain, Kartik Gupta
- #92 Globally Optimal Gradient Descent for a ConvNet with Gaussian Inputs**
Alon Brutzkus, Amir Globerson
- #93 Re-revisiting Learning on Hypergraphs: Confidence Interval and Subgradient Method**
Chenzi Zhang, Shuguang Hu, Zhihao Gavin Tang, Hubert Chan
- #94 Meta Networks**
Tsendsuren Munkhdalai, Hong Yu
- #95 Bottleneck Conditional Density Estimation**
Rui Shu, Hung Bui, Mohammad Ghavamzadeh
- #96 Exploiting Strong Convexity from Data with Primal-Dual First-Order Algorithms**
Jialei Wang, Lin Xiao
- #97 Interactive Learning from Policy-Dependent Human Feedback**
James MacGlashan, Mark Ho, Robert Loftin, Bei Peng, Guan Wang, David L Roberts, Matthew E. Taylor, Michael L. Littman
- #98 Learning to Discover Sparse Graphical Models**
Eugene Belilovsky, Kyle Kastner, Gael Varoquaux, Matthew B Blaschko
- #99 On Context-Dependent Clustering of Bandits**
Claudio Gentile, Shuai Li, Puru Kar, Alexandros Karatzoglou, Giovanni Zappella, Evans Etrue Howard
- #100 Provable Alternating Gradient Descent for Non-negative Matrix Factorization with Strong Correlations**
Yuanzhi Li, Yingyu Liang
- #101 Convexified Convolutional Neural Networks**
Yuchen Zhang, Percy Liang, Martin Wainwright
- #102 Self-Paced Co-training**
Fan Ma, Deyu Meng, Qi Xie, Zina Li, Xuanyi Dong
- #103 SplitNet: Learning to Semantically Split Deep Networks for Parameter Reduction and Model Parallelization**
Juyong Kim, Yookoon Park, Gunhee Kim, Sung Ju Hwang
- #104 Learning Deep Latent Gaussian Models with Markov Chain Monte Carlo**
Matthew Hoffman
- #105 Doubly Greedy Primal-Dual Coordinate Descent for Sparse Empirical Risk Minimization**
Qi Lei, Ian Yen, Chao-Yuan Wu, Inderjit Dhillon, Pradeep Ravikumar
- #106 End-to-End Differentiable Adversarial Imitation Learning**
Nir Baram, Oron Anshel, Itai Caspi, Shie Mannor
- #107 Local-to-Global Bayesian Network Structure Learning**
Tian Gao, Kshitij Fadnis, Murray Campbell
- #108 Provably Optimal Algorithms for Generalized Linear Contextual Bandits**
Lihong Li, Yu Lu, Dengyong Zhou
- #109 No Spurious Local Minima in Nonconvex Low Rank Problems: A Unified Geometric Analysis**
Rong Ge, Chi Jin, Yi Zheng



- #110 On the Expressive Power of Deep Neural Networks**
Maithra Raghu, Ben Poole, Surya Ganguli, Jon Kleinberg, Jascha Sohl-Dickstein
- #111 Semi-Supervised Classification Based on Classification from Positive and Unlabeled Data**
Tomoya Sakai, Marthinus C du Plessis, Gang Niu, Masashi Sugiyama
- #112 Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks**
Chelsea Finn, Pieter Abbeel, Sergey Levine
- #113 Zero-Inflated Exponential Family Embeddings**
Liping Liu, David Blei
- #114 A Richer Theory of Convex Constrained Optimization with Reduced Projections and Improved Rates**
Tianbao Yang, Qihang Lin, Lijun Zhang
- #115 Learning in POMDPs with Monte Carlo Tree Search**
Sammie Katt, Frans A Oliehoek, Chris Amato
- #116 Composing Tree Graphical Models with Persistent Homology Features for Clustering Mixed-Type Data**
XIUYAN NI, Novi Quadrianto, Yusu Wang, Chao Chen
- #117 Safety-Aware Algorithms for Adversarial Contextual Bandit**
Wen Sun, Debadeepta Dey, Ashish Kapoor
- #118 Coherence Pursuit: Fast, Simple, and Robust Subspace Recovery**
Mostafa Rahmani, George Atia
- #119 Depth-Width Tradeoffs in Approximating Natural Functions With Neural Networks**
Itay Safran, Ohad Shamir
- #120 Iterative Machine Teaching**
Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda Smith, Jim Rehg, Le Song
- #121 AdaNet: Adaptive Structural Learning of Artificial Neural Networks**
Corinna Cortes, Xavi Gonzalvo, Vitaly Kuznetsov, Mehryar Mohri, Scott Yang
- #122 Convex Phase Retrieval without Lifting via PhaseMax**
Tom Goldstein, Christoph Studer
- #123 DARLA: Improving Zero-Shot Transfer in Reinforcement Learning**
Irina Higgins, Arka Pal, Andrei A Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, Alexander Lerchner
- #124 On Relaxing Determinism in Arithmetic Circuits**
Arthur Choi, Adnan Darwiche
- #125 Adaptive Multiple-Arm Identification**
Jiecao (Jack) Chen, Xi Chen, Qin Zhang, Yuan Zhou
- #126 Tensor Decomposition with Smoothness**
Masaaki Imaizumi, Kohei Hayashi
- #127 Automated Curriculum Learning for Neural Networks**
Alex Graves, Marc Bellemare, Jacob Menick, Remi Munos, Koray Kavukcuoglu
- #128 Attentive Recurrent Comparators**
Pranav Shyam, Shubham Gupta, Ambedkar Dukkipati
- #129 An Infinite Hidden Markov Model With Similarity-Biased Transitions**
Colin Dawson, Bill Huang, Clayton T. Morrison
- #130 Efficient Nonmyopic Active Search**
Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, Roman Garnett
- #131 Asymmetric Tri-training for Unsupervised Domain Adaptation**
Saito Kuniaki, Yoshitaka Ushiku, Tatsuya Harada
- #132 State-Frequency Memory Recurrent Neural Networks**
Hao Hu, Guo-Jun Qi
- #133 Batched High-dimensional Bayesian Optimization via Structural Kernel Learning**
Zi Wang, Chengtao Li, Stefanie Jegelka, Pushmeet Kohli
- #134 Leveraging Union of Subspace Structure to Improve Constrained Clustering**
John Lipor, Laura Balzano
- #135 Source-Target Similarity Modelings for Multi-Source Transfer Gaussian Process Regression**
Pengfei Wei, Ramon Sagarna, Yiping Ke, Yew Ong, Chi Goh
- #136 Delta Networks for Optimized Recurrent Network Computation**
Daniel Neil, Jun Lee, Tobi Delbruck, Shih-Chii Liu
- #137 From Patches to Images: A Nonparametric Generative Model**
Geng Ji, Michael C. Hughes, Erik Sudderth
- #138 Active Heteroscedastic Regression**
Kamalika Chaudhuri, Prateek Jain, Nagarajan Natarajan
- #139 Multi-task Learning with Labeled and Unlabeled Tasks**
Anastasia Pentina, Christoph Lampert
- #140 Recurrent Highway Networks**
Julian Zilly, Rupesh Srivastava, Jan Koutnik, Jürgen Schmidhuber
- #141 Fast Bayesian Intensity Estimation for the Permenental Process**
Christian Walder, Adrian N Bishop
- #142 Active Learning for Cost-Sensitive Classification**
Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, Hal Daumé III, John Langford
- #143 Schema Networks: Zero-shot Transfer with a Generative Causal Model of Intuitive Physics**
Ken Kansky, Tom Silver, David A Mély, Mo Eldawy, Miguel Lazaro-Gredilla, Xinghua Lou, Nimrod Dorfman, Szymon Sidor, Scott Phoenix, Dileep George
- #144 A Birth-Death Process for Feature Allocation**
Konstantina Palla, David Knowles, Zoubin Ghahramani
- #145 Diameter-Based Active Learning**
Chris Tosh, Sanjoy Dasgupta
- #146 Risk Bounds for Transferring Representations With and Without Fine-Tuning**
Daniel McNamara, Nina Balcan

Tuesday

AUGUST 8TH | SESSIONS



TIME	DESCRIPTION	LOCATION	TIME	DESCRIPTION	LOCATION
9 am	Test Of Time Award	L2, Darling Harbour	3 pm	Break	
10 am	Break		3:30 pm	SESSION THREE	
10:30 am	SESSION ONE			Deep learning 7: Analysis	L2, Darling Harbour
	Deep learning 5:			Recurrent neural networks 4	L2, Parkside 1
	Fisher approximations	L2, Darling Harbour		Continuous optimization 7	L2, Parkside 2
	Recurrent neural networks 2	L2, Parkside 1		Game theory and multiagents	L4, C4.5
	Continuous optimization 5	L2, Parkside 2		Probabilistic inference 3	L4, C4.9 & C4.10
	Reinforcement learning 5	L4, C4.5		ML and programming	L4, C4.1
	Probabilistic inference 1	L4, C4.9 & C4.10		Sparsity 2	L4, C4.4
	Networks & relational learning	L4, C4.1		Information theory	L4, C4.8
	Spectral methods	L4, C4.4		Metric learning	L4, C4.6 & C4.7
	Learning theory 1	L4, C4.8	5:15 pm	Invited Talk: Peter Donnelly	
	Clustering 1	L4, C4.6 & C4.7		Genomics, Big Data, and Machine Learning: Understanding the Human Wiring Diagram and Driving the Healthcare Revolution	L2, Darling Harbour
12 pm	LUNCH (On Your Own)		6:15 pm	Poster Session	L2, Gallery
1:30 pm	SESSION TWO				
	Deep learning 6	L2, Darling Harbour			
	Recurrent neural networks 3	L2, Parkside 1			
	Continuous optimization 6	L2, Parkside 2			
	Ensemble methods	L4, C4.5			
	Probabilistic inference 2	L4, C4.9 & C4.10			
	Ranking and preferences	L4, C4.1			
	Sparsity 1	L4, C4.4			
	Learning theory 2	L4, C4.8			
	Clustering 2	L4, C4.6 & C4.7			



SESSION ONE - 10:30 AM - 12:00 PM

Deep learning 5: Fisher approximations

Location: Level 2, Darling Harbour

- **Relative Fisher Information and Natural Gradient for Learning Large Modular Models**
Ke Sun, Frank Nielsen
- **Learning Deep Architectures via Generalized Whitened Neural Networks**
Ping Luo
- **Continual Learning Through Synaptic Intelligence**
Friedemann Zenke, Ben Poole, Surya Ganguli
- **Adaptive Neural Networks for Efficient Inference**
Tolga Bolukbasi, Joe Wang, Ofer Dekel, Venkatesh Saligrama
- **Combined Group and Exclusive Sparsity for Deep Neural Networks**
jaehong yoon, Sung Ju Hwang

Recurrent neural networks 2

Location: Level 2, Parkside 1

- **Efficient Orthogonal Parametrisation of Recurrent Neural Networks Using Householder Reflections**
Zakaria mhammedi, Andrew Hellicar, James Bailey, Ashfaqur Rahman
- **On orthogonality and learning RNNs with long term dependencies**
Eugene Vorontsov, Chiheb Trabelsi, Christopher Pal, Samuel Kadoury
- **Tunable Efficient Unitary Neural Networks (EUNN) and their application to RNNs**
Li Jing, Yichen Shen, Tena Dubcek, John E Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, Marin Solja{\v{c}}i{\{c}}
- **The Statistical Recurrent Unit**
Junier Oliva, Barnabás Póczos, Jeff Schneider
- **Input Switched Affine Networks: An RNN Architecture Designed for Interpretability**
Jakob Foerster, Justin Gilmer, Jan Chorowski, Jascha Sohl-Dickstein, David Sussillo

Continuous optimization 5

Location: Level 2, Parkside 2

- **Lazifying Conditional Gradient Algorithms**
Gábor Braun, Sebastian Pokutta, Daniel Zink
- **Conditional Accelerated Lazy Stochastic Gradient Descent**
Guanghui, Sebastian Pokutta, Yi Zhou, Daniel Zink
- **SARAH: A Novel Method for Machine Learning Problems Using Stochastic Recursive Gradient**
MLTD Nguyen, Jie Liu, Katya Scheinberg, Martin Takac
- **Approximate Steepest Coordinate Descent**
Sebastian Stich, Anant Raj, Martin Jaggi
- **StingyCD: Safely Avoiding Wasteful Updates in Coordinate Descent**
Tyler Johnson, Carlos Guestrin

Reinforcement learning 5

Location: Level 4, C4.5

- **Data-Efficient Policy Evaluation Through Behavior Policy Search**
Josiah Hanna, Philip S. Thomas, Peter Stone, Scott Niekum
- **Stochastic Variance Reduction Methods for Policy Evaluation**
Simon Du, Jianshu Chen, Lihong Li, Lin Xiao, Dengyong Zhou
- **Optimal and Adaptive Off-policy Evaluation in Contextual Bandits**
Yu-Xiang Wang, Alekh Agarwal, Miroslav Dudik
- **Consistent On-Line Off-Policy Evaluation**
Assaf Hallak, Shie Mannor
- **Contextual Decision Processes with low Bellman rank are PAC-Learnable**
Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert Schapire



SESSION ONE - 10:30 AM - 12:00 PM

Probabilistic inference 1

Location: Level 4, C4.9 & C4.10

- **Exact MAP Inference by Avoiding Fractional Vertices**
Erik Lindgren, Alex Dimakis, Adam Klivans
- **Exact Inference for Integer Latent-Variable Models**
Kevin Winner, Debora Sujono, Daniel Sheldon
- **Improving Viterbi is Hard: Better Runtimes Imply Faster Clique Algorithms**
Arturs Backurs, Christos Tzamos
- **Variational Inference for Sparse and Undirected Models**
John Ingraham, Debora Marks
- **Tensor Belief Propagation**
Andrew Wrigley, Wee Sun Lee Lee, Nan Ye

Networks and relational learning

Location: Level 4, C4.1

- **Leveraging Node Attributes for Incomplete Relational Data**
He Zhao, Lan Du Du, Wray Buntine
- **Bayesian inference on random simple graphs with power law degree distributions**
Juho Lee, Creighton Heaukulani, Zoubin Ghahramani, Lancelot F. James, Seungjin Choi
- **Analogical Inference for Multi-relational Embeddings**
Hanxiao Liu, Yuexin Wu, Yiming Yang
- **Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs**
Rakshit Trivedi, Hajun Dai, Yichen Wang, Le Song
- **Deep Generative Models for Relational Data with Side Information**
Changwei Hu, Piyush Rai, Lawrence Carin

Spectral methods

Location: Level 4, C4.4

- **How Close Are the Eigenvectors of the Sample and Actual Covariance Matrices?**
Andreas Loukas
- **Faster Principal Component Regression and Stable Matrix Chebyshev Approximation**
Zeyuan Allen-Zhu, Yuanzhi Li

- **Spectral Learning from a Single Trajectory under Finite-State Policies**

Borja de Balle Pigem, Odalric Maillard

- **Capacity Releasing Diffusion for Speed and Locality.**
Di Wang, Kimon Fountoulakis, Monika Henzinger, Michael Mahoney, Satish Rao
- **Doubly Accelerated Methods for Faster CCA and Generalized Eigendecomposition**
Zeyuan Allen-Zhu, Yuanzhi Li

Learning theory 1

Location: Level 4, C4.8

- **Sketched Ridge Regression: Optimization Perspective, Statistical Perspective, and Model Averaging**
Shusen Wang, Alex Gittens, Michael Mahoney
- **Estimating the unseen from multiple populations**
Aditi Raghunathan, Greg Valiant, James Zou
- **Meritocratic Fairness for Cross-Population Selection**
Michael Kearns, Aaron Roth, Steven Wu
- **Neural networks and rational functions**
Matus Telgarsky

Clustering 1

Location: Level 4, C4.6 & C4.7

- **Distributed and Provably Good Seedings for k-Means in Constant Rounds**
Olivier Bachem, Mario Lucic, Andreas Krause
- **Consistent k-Clustering**
Silvio Lattanzi, Sergei Vassilvitskii
- **Towards K-means-friendly Spaces: Simultaneous Deep Learning and Clustering**
Bo Yang, Xiao Fu, Nicholas Sidiropoulos, Mingyi Hong
- **Hyperplane Clustering Via Dual Principal Component Pursuit**
Manolis Tsakiris, Rene Vidal
- **Multilevel Clustering via Wasserstein Means**
Nhat Ho, Long Nguyen, Mikhail Yurochkin, Hung Bui, Viet Huynh, Dinh Phung



SESSION TWO - 1:30 PM - 3:00 PM

Deep learning 6

Location: Level 2, Darling Harbour

- **Input Convex Neural Networks**
Brandon Amos, Lei Xu, Zico Kolter
- **OptNet: Differentiable Optimization as a Layer in Neural Networks**
Brandon Amos, Zico Kolter
- **Parseval Networks: Improving Robustness to Adversarial Examples**
Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, Nicolas Usunier
- **Regularising Non-linear Models Using Feature Side-information**
Amina Mollaysa, Pablo Strasser, Alexandros Kalousis

Recurrent neural networks 3

Location: Level 2, Parkside 1

- **Online and Linear-Time Attention by Enforcing Monotonic Alignments**
Colin Raffel, Thang Luong, Peter Liu, Ron Weiss, Douglas Eck
- **Sequence Tutor: Conservative fine-tuning of sequence generation models with KL-control**
Natasha Jaques, Shixiang Gu, Dzmitry Bahdanau, Jose Hernandez-Lobato, Richard E Turner, Douglas Eck
- **Deep Voice: Real-time Neural Text-to-Speech**
agibiansky Gibiansky, Mike Chrzanowski, Mohammad Shoeybi, Shubho Sengupta, Gregory Diamos, Serkan Arik, Jonathan Raiman, John Miller, Xian Li, Yongguo Kang, Adam Coates, Andrew Ng
- **DeepBach: a Steerable Model for Bach Chorales Generation**
Gaëtan HADJERES, François Pachet, Frank Nielsen
- **Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders**
Cinjon Resnick, Adam Roberts, JesseEngel Engel, Douglas Eck, Sander Dieleman, Karen Simonyan, Mohammad Norouzi

Continuous optimization 6

Location: Level 2, Parkside 2

- **Stochastic modified equations and adaptive stochastic gradient algorithms**
Qianxiao Li, Cheng Tai, Weinan E

- **Dissipativity Theory for Nesterov's Accelerated Method**

Bin Hu, Laurent Lessard

- **An Analytical Formula of Population Gradient for two-layered ReLU network and its Applications in Convergence and Critical Point Analysis**

Yuandong Tian

- **Forward and Reverse Gradient-Based Hyperparameter Optimization**

Luca Franceschi, Michele Donini, Paolo Frasconi, Massimiliano Pontil

- **Adaptive Sampling Probabilities for Non-Smooth Optimization**

Hongseok Namkoong, Aman Sinha, Steve Yadlowsky, John Duchi

Ensemble methods

Location: Level 4, C4.5

- **A Simple Multi-Class Boosting Framework with Theoretical Guarantees and Empirical Proficiency**
Ron Appel, Pietro Perona
- **Gradient Boosted Decision Trees for High Dimensional Sparse Output**
Si Si, Huan Zhang, Sathiya Keerthi, Dhruv Mahajan, Inderjit Dhillon, Cho-Jui Hsieh
- **Globally Induced Forest: A Prepruning Compression Scheme**
Jean-Michel Begon, Arnaud Joly, Pierre Geurts
- **Forest-type Regression with General Losses and Robust Forest**
Hanbo Li, Andy Martin
- **Confident Multiple Choice Learning**
Kimin Lee, Changho Hwang, Kyoungsoo Park, Jinwoo Shin

Probabilistic inference 2

Location: Level 4, C4.9& C4.10

- **Faster Greedy MAP Inference for Determinantal Point Processes**

Insu Han, Prabhanjan Kambadur, Kyoungsoo Park, Jinwoo Shin

- **Zonotope hit-and-run for efficient sampling from projection DPPs**

Guillaume Gautier, Rémi Bardenet, Michal Valko



SESSION TWO - 1:30 PM - 3:00 PM

- **A Divergence Bound for Hybrids of MCMC and Variational Inference and an Application to Langevin Dynamics and SGVI**

Justin Domke

- **On the Sampling Problem for Kernel Quadrature**

Francois-Xavier Briol, Chris J Oates, Jon Cockayne, Wilson Ye Chen, Mark Girolami

- **Measuring Sample Quality with Kernels**

Jackson Gorham, Lester Mackey

Ranking and preferences

Location: Level 4, C4.1

- **ChoiceRank: Identifying Preferences from Node Traffic in Networks**

Ium Maystre, Matt Grossglauser

- **Statistical Inference for Incomplete Ranking Data: The Case of Rank-Dependent Coarsening**

Mohsen Ahmadi Fahandar, Eyke Hüllermeier, Ines Couso

- **Just Sort It! A Simple and Effective Approach to Active Preference Learning**

Ium Maystre, Matt Grossglauser

- **Maximum Selection and Ranking under Noisy Comparisons**

Moein Falahatgar, Alon Orlitsky, Venkatadheeraj Pichapati, Ananda Suresh

- **Active Learning for Top- k Rank Aggregation from Noisy Comparisons**

Soheil Mohajer, Changho Suh, Adel Elmahdy

Sparsity 1

Location: Level 4, C4.4

- **On the Iteration Complexity of Support Recovery via Hard Thresholding Pursuit**

Jie Shen, Ping Li

- **Dual Iterative Hard Thresholding: From Non-convex Sparse Minimization to Non-smooth Concave Maximization**

Bo Liu, Xiaotong Yuan, Lezi Wang, Qingshan Liu, Dimitris Metaxas

- **On The Projection Operator to A Three-view Cardinality Constrained Set**

Haichuan Yang, Shupeng Gui, Chuyang Ke, Daniel Stefankovic, Ryohei Fujimaki, Ji Liu

- **Sparse + Group-Sparse Dirty Models: Statistical Guarantees without Unreasonable Conditions and a Case for Non-Convexity**

Eunho Yang, Aurelie Lozano

- **Compressed Sensing using Generative Models**

Ashish Bora, Ajil Jalal, Eric Price, Alex Dimakis

Learning theory 2

Location: Level 4, C4.8

- **Uniform Deviation Bounds for k-Means Clustering**

Olivier Bachem, Mario Lucic, Hamed Hassani, Andreas Krause

- **Uniform Convergence Rates for Kernel Density Estimation**

Heinrich Jiang

- **Density Level Set Estimation on Manifolds with DBSCAN**

Heinrich Jiang

- **Algorithmic Stability and Hypothesis Complexity**

Tongliang Liu, Gábor Lugosi, Gergely Neu, Dacheng Tao

- **Consistency Analysis for Binary Classification Revisited**

Krzysztof Dembczynski, Wojciech Kotłowski, Sanmi Koyejo, Nagarajan Natarajan

Clustering 2

Location: Level 4, C4.6 & C4.7

- **Co-clustering through Optimal Transport**

Charlotte Laclau, Ievgen Redko, Basarab Matei, Younès Bennani, Vincent Brault

- **Multiple Clustering Views from Multiple Uncertain Experts**

Yale Chang, Junxiang Chen, Michael Cho, Peter Castaldi, Edwin Silverman, Jennifer G Dy

- **Clustering by Sum of Norms: Stochastic Incremental Algorithm, Convergence and Cluster Recovery**

Ashkan Panahi, Devdatt Dubhashi, Fredrik D Johansson, Chiranjib Bhattacharya

- **Clustering High Dimensional Dynamic Data Streams**

Lin Yang, Harry Lang, Christian Sohler, Vladimir Braverman, Gereon Frahling



SESSION THREE - 3:30 PM - 5:00 PM

Deep learning 7: analysis

Location: Level 2, Darling Harbour

- **A Closer Look at Memorization in Deep Networks**
David Krueger, Yoshua Bengio, Stanislaw Jastrzebski, Maxinder S. Kanwal, Nicolas Ballas, Asja Fischer, Emmanuel Bengio, Devansh Arpit, Tegan Maharaj, Aaron Courville, Simon Lacoste-Julien
- **Cognitive Psychology for Deep Neural Networks: A Shape Bias Case Study**
Sam Ritter, David GT Barrett, Adam Santoro, Matthew Botvinick
- **Visualizing and Understanding Multilayer Perceptron Models: A Case Study in Speech Processing**
Tasha Nagamine, Nima Mesgarani
- **Axiomatic Attribution for Deep Networks**
Mukund Sundarajan, Ankur Taly, Qi Qi Yan
- **On Calibration of Modern Neural Networks**
Chuan Guo, Geoff Pleiss, Yu Sun, Kilian Weinberger

Recurrent neural networks 4

Location: Level 2, Parkside 1

- **Learning to Generate Long-term Future via Hierarchical Prediction**
Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, Honglak Lee
- **Sequence to Better Sequence: Continuous Revision of Combinatorial Structures**
Jonas Mueller, David Gifford, Tommi Jaakkola
- **Tensor-Train Recurrent Neural Networks for Video Classification**
Yinchong Yang, Denis Kropas, Volker Tresp
- **Sequence Modeling via Segmentations**
Chong Wang, Yining Wang, Po-Sen Huang, Abdelrahman Mohammad, Dengyong Zhou, Li Deng
- **Latent LSTM Allocation: Joint clustering and non-linear dynamic modeling of sequence data**
Manzil Zaheer, Amr Ahmed, Alex Smola

Continuous optimization 7

Location: Level 2, Parkside 2

- **Sub-sampled Cubic Regularization for Non-convex Optimization**
Jonas Kohler, Aurelien Lucchi

- **Natasha: Faster Non-Convex Stochastic Optimization Via Strongly Non-Convex Parameter**
Zeyuan Allen-Zhu
- **“Convex Until Proven Guilty”: Dimension-Free Acceleration of Gradient Descent on Non-Convex Functions**
Yair Carmon, John Duchi, Oliver Hinder, Aaron Sidford
- **Convergence Analysis of Proximal Gradient with Momentum for Nonconvex Optimization**
Qunwei Li, Yi Zhou, Yingbin Liang, Pramod K Varshney
- **How to Escape Saddle Points Efficiently**
Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M. Kakade, Michael Jordan

Game theory and multiagents

Location: Level 4, C4.5

- **Regret Minimization in Behaviorally-Constrained Zero-Sum Games**
Gabriele Farina, Christian Kroer, Tuomas Sandholm
- **Reduced Space and Faster Convergence in Imperfect-Information Games via Pruning**
Noam Brown, Tuomas Sandholm
- **Strongly-Typed Agents are Guaranteed to Interact Safely**
David Balduzzi
- **Coordinated Multi-Agent Imitation Learning**
Hoang Le, Yisong Yue, Peter Carr, Patrick Lucey
- **Deep Decentralized Multi-task Multi-Agent Reinforcement Learning under Partial Observability**
Shayegan Omidshafiei, Jason Pazis, Chris Amato, Jonathan How, John L Vian

Probabilistic inference 3

Location: Level 4, C4.9& C4.10

- **Variational Boosting: Iteratively Refining Posterior Approximations**
Andrew Miller, Nick J Foti, Ryan Adams
- **Lost Relatives of the Gumbel Trick**
Matej Balog, Nilesh Tripuraneni, Zoubin Ghahramani, Adrian Weller
- **Learning to Aggregate Ordinal Labels by Maximizing Separating Width**
Guangyong Chen, Shengyu Zhang, Di Lin, HUI Huang, Pheng Heng



SESSION THREE - 3:30 PM - 5:00 PM

- **Uncorrelation and Evenness: a New Diversity-Promoting Regularizer**

Pengtao Xie, Aarti Singh, Eric Xing

- **Learning Latent Space Models with Angular Constraints**

Pengtao Xie, Yuntian Deng, Yi Zhou, Abhimanu Kumar, Yaoliang Yu, James Zou, Eric Xing

ML and programming

Location: Level 4, C4.1

- **Learning to Align the Source Code to the Compiled Object Code**

Ariel Green, Lior Wolf

- **RobustFill: Neural Program Learning under Noisy I/O**

Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdelrahman Mohammad, Pushmeet Kohli

- **Programming with a Differentiable Forth Interpreter**

Matko Bošnjak, Tim Rocktäschel, Jason Naradowsky, Sebastian Riedel

- **Differentiable Programs with Neural Libraries**

Alex Gaunt, Marc Brockschmidt, Nate Kushman, Daniel Tarlow

- **Developing Bug-Free Machine Learning Systems With Formal Mathematics**

Daniel Selsam, Percy Liang, David L Dill

Sparsity 2

Location: Level 4, C4.4

- **Scaling Up Sparse Support Vector Machines by Simultaneous Feature and Sample Reduction**

Weizhong Zhang, Bin Hong, Wei Liu, Jieping Ye, Deng Cai, Xiaofei He, Jie Wang

- **Efficient Distributed Learning with Sparsity**

Jialei Wang, Mladen Kolar, Nati Srebro, Tong Zhang

- **Innovation Pursuit: A New Approach to the Subspace Clustering Problem**

Mostafa Rahmani, George Atia

- **Selective Inference for Sparse High-Order Interaction Models**

Shinya Suzumura, Kazuya Nakagawa, Yuta Umezumi, Koji Tsuda, Ichiro Takeuchi

- **Dictionary Learning Based on Sparse Distribution Tomography**

Pedram Pad, Farnood Salehi, Elisa Celis, Patrick Thiran, Michael Unser

Information theory

Location: Level 4, C4.8

- **Distributed Mean Estimation with Limited Communication**

Ananda Suresh, Felix Yu, Sanjiv Kumar, Brendan McMahan

- **Nonparanormal Information Estimation**

Shashank Singh, Barnabás Póczos

- **A Unified Maximum Likelihood Approach for Estimating Symmetric Properties of Discrete Distributions**

Jayadev Acharya, Hirakendu Das, Alon Orlitsky, Ananda Suresh

- **Gradient Coding: Avoiding Stragglers in Distributed Learning**

Rashish Tandon, Qi Lei, Alex Dimakis, NIKOS KARAMPATZIAKIS

- **Learning Discrete Representations via Information Maximizing Self-Augmented Training**

Weihua Hu, Takeru Miyato, Seiya Tokui, Eiichi Matsumoto, Masashi Sugiyama

Metric learning

Location: Level 4, C4.6 & C4.7

- **Fast k-Nearest Neighbour Search via Prioritized DCI**

Ke Li, Jitendra Malik

- **Deep Spectral Clustering Learning**

Marc Law, Raquel Urtasun, Zemel Rich

- **Joint Dimensionality Reduction and Metric Learning: A Geometric Take**

Mehrtash Harandi, Mathieu Salzmann, Richard I Hartley

- **ProtoNN: Compressed and Accurate kNN for Resource-scarce Devices**

Chirag Gupta, ARUN SUGGALA, Ankit Goyal, Saurabh Goyal, Ashish Kumar, Bhargavi Paranjape, Harsha Vardhan Simhadri, Raghavendra Udupa, Manik Varma, Prateek Jain



- #1 **The loss surface of deep and wide neural networks**
Quynh Nguyen, Matthias Hein
- #2 **Neural Taylor Approximations: Convergence and Exploration in Rectifier Networks**
David Balduzzi, Brian McWilliams, Tony Butler-Yeoman
- #3 **Sharp Minima Can Generalize For Deep Nets**
Laurent Dinh, Razvan Pascanu, Samy Bengio, Yoshua Bengio
- #4 **Geometry of Neural Network Loss Surfaces via Random Matrix Theory**
Jeffrey Pennington, Yasaman Bahri
- #5 **The Shattered Gradients Problem: If resnets are the answer, then what is the question?**
David Balduzzi, Marcus Frean, Wan-Duo Ma, Brian McWilliams, Lennox Leary, J.P. Lewis
- #6 **Learning to Learn without Gradient Descent by Gradient Descent**
Yutian Chen, Matthew Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Tim Lillicrap, Matthew Botvinick, Nando de Freitas
- #7 **A Semismooth Newton Method for Fast, Generic Convex Programming**
Alnur Ali, Eric Wong, Zico Kolter
- #8 **Unifying task specification in reinforcement learning**
Martha White
- #9 **Efficient Online Bandit Multiclass Learning with $O(\sqrt{T})$ Regret**
Alina Beygelzimer, Francesco Orabona, Chicheng Zhang
- #10 **Orthogonalized ALS: A Theoretically Principled Tensor Decomposition Algorithm for Practical Use**
Vatsal Sharan, Gregory Valiant
- #11 **Learned Optimizers that Scale and Generalize**
Olga Wichrowska, Niru Maheswaranathan, Matthew Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Nando de Freitas, Jascha Sohl-Dickstein
- #12 **Approximate Newton Methods and Their Local Convergence**
Haishan Ye, Luo Luo, Zhihua Zhang
- #13 **A Distributional Perspective on Reinforcement Learning**
Marc Bellemare, Will Dabney, Remi Munos
- #14 **Active Learning for Accurate Estimation of Linear Models**
Carlos Riquelme Ruiz, Mohammad Ghavamzadeh, Alessandro Lazaric
- #15 **Tensor Decomposition via Simultaneous Power Iteration**
Poan Wang, Chi-Jen Lu
- #16 **Learning Gradient Descent: Better Generalization and Longer Horizons**
Kaifeng Lv, Shunhua Jiang, Jian Li
- #17 **Stochastic Adaptive Quasi-Newton Methods for Minimizing Expected Values**
Chaoxu Zhou, Wenbo Gao, Donald Goldfarb
- #18 **Hierarchy Through Composition with Multitask LMDPs**
Andrew Saxe, Adam Earle, Benjamin Rosman
- #19 **Adaptive Feature Selection: Computationally Efficient Online Sparse Linear Regression under RIP**
Satyen Kale, Zohar Karnin, Tengyuan Liang, David Pal
- #20 **A Unified Variance Reduction-Based Framework for Nonconvex Low-Rank Matrix Recovery**
Lingxiao Wang, Xiao Zhang, Quanquan Gu
- #21 **Learning Algorithms for Active Learning**
Philip Bachman, Alessandro Sordoni, Adam Trischler
- #22 **Practical Gauss-Newton Optimisation for Deep Learning**
Alex Botev, Hippolyt Ritter, David Barber
- #23 **A Laplacian Framework for Option Discovery in Reinforcement Learning**
Marlos C. Machado, Marc Bellemare, Michael Bowling
- #24 **Emulating the Expert: Inverse Optimization through Online Learning**
Sebastian Pokutta, Andreas Bärmann, Oskar Schneider
- #25 **An Efficient, Sparsity-Preserving, Online Algorithm for Low-Rank Approximation**
Dave Anderson, Ming Gu
- #26 **Tensor Balancing on Statistical Manifold**
Mahito Sugiyama, Hiroyuki Nakahara, Koji Tsuda
- #27 **Modular Multitask Reinforcement Learning with Policy Sketches**
Jacob Andreas, Dan Klein, Sergey Levine
- #28 **Variants of RMSProp and Adagrad with Logarithmic Regret Bounds**
Mahesh Chandra Mukkamala, Matthias Hein
- #29 **Algorithms for ℓ_p Low-Rank Approximation**
Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, David Woodruff
- #30 **Relative Fisher Information and Natural Gradient for Learning Large Modular Models**
Ke Sun, Frank Nielsen
- #31 **Efficient Orthogonal Parametrisation of Recurrent Neural Networks Using Householder Reflections**
Zakaria mhammedi, Andrew Helicar, James Bailey, Ashfaqur Rahman
- #32 **Lazifying Conditional Gradient Algorithms**
Gábor Braun, Sebastian Pokutta, Daniel Zink
- #33 **Data-Efficient Policy Evaluation Through Behavior Policy Search**
Josiah Hanna, Philip S. Thomas, Peter Stone, Scott Niekum
- #34 **Exact MAP Inference by Avoiding Fractional Vertices**
Erik Lindgren, Alex Dimakis, Adam Klivans
- #35 **Leveraging Node Attributes for Incomplete Relational Data**
He Zhao, Lan Du Du, Wray Buntine
- #36 **How Close Are the Eigenvectors of the Sample and Actual Covariance Matrices?**
Andreas Loukas

- #37 Distributed and Provably Good Seedings for k-Means in Constant Rounds**
Olivier Bachem, Mario Lucic, Andreas Krause
- #38 Learning Deep Architectures via Generalized Whitened Neural Networks**
Ping Luo
- #39 On orthogonality and learning RNNs with long term dependencies**
Eugene Vorontsov, Chiheb Trabelsi, Christopher Pal, Samuel Kadoury
- #40 Conditional Accelerated Lazy Stochastic Gradient Descent**
Guanghui, Sebastian Pokutta, Yi Zhou, Daniel Zink
- #41 Stochastic Variance Reduction Methods for Policy Evaluation**
Simon Du, Jianshu Chen, Lihong Li, Lin Xiao, Dengyong Zhou
- #42 Exact Inference for Integer Latent-Variable Models**
Kevin Winner, Debora Sujono, Daniel Sheldon
- #43 Bayesian inference on random simple graphs with power law degree distributions**
Juho Lee, Creighton Heaukulani, Zoubin Ghahramani, Lancelot F. James, Seungjin Choi
- #44 Faster Principal Component Regression and Stable Matrix Chebyshev Approximation**
Zeyuan Allen-Zhu, Yuanzhi Li
- #45 Consistent k-Clustering**
Silvio Lattanzi, Sergei Vassilvitskii
- #46 Continual Learning Through Synaptic Intelligence**
Friedemann Zenke, Ben Poole, Surya Ganguli
- #47 Tunable Efficient Unitary Neural Networks (EUNN) and their application to RNNs**
Li Jing, Yichen Shen, Tena Dubcek, John E Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, Marin Soljačić
- #48 SARAH: A Novel Method for Machine Learning Problems Using Stochastic Recursive Gradient**
MLTD Nguyen, Jie Liu, Katya Scheinberg, Martin Takac
- #49 Optimal and Adaptive Off-policy Evaluation in Contextual Bandits**
Yu-Xiang Wang, Alekh Agarwal, Miroslav Dudik
- #50 Improving Viterbi is Hard: Better Runtimes Imply Faster Clique Algorithms**
Arturs Backurs, Christos Tzamos
- #51 Analogical Inference for Multi-relational Embeddings**
Hanxiao Liu, Yuexin Wu, Yiming Yang
- #52 Spectral Learning from a Single Trajectory under Finite-State Policies**
Borja de Balle Pigem, Odalric Maillard
- #53 Towards K-means-friendly Spaces: Simultaneous Deep Learning and Clustering**
Bo Yang, Xiao Fu, Nicholas Sidiropoulos, Mingyi Hong
- #54 Adaptive Neural Networks for Efficient Inference**
Tolga Bolukbasi, Joe Wang, Ofer Dekel, Venkatesh Saligrama
- #55 The Statistical Recurrent Unit**
Junier Oliva, Barnabás Póczos, Jeff Schneider
- #56 Approximate Steepest Coordinate Descent**
Sebastian Stich, Anant Raj, Martin Jaggi
- #57 Consistent On-Line Off-Policy Evaluation**
Assaf Hallak, Shie Mannor
- #58 Variational Inference for Sparse and Undirected Models**
John Inghram, Debora Marks
- #59 Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs**
Rakshit Trivedi, Hajun Dai, Yichen Wang, Le Song
- #60 Capacity Releasing Diffusion for Speed and Locality.**
Di Wang, Kimon Fountoulakis, Monika Henzinger, Michael Mahoney, Satish Rao
- #61 Hyperplane Clustering Via Dual Principal Component Pursuit**
Manolis Tsakiris, Rene Vidal
- #62 Combined Group and Exclusive Sparsity for Deep Neural Networks**
jaehong yoon, Sung Ju Hwang
- #63 Input Switched Affine Networks: An RNN Architecture Designed for Interpretability**
Jakob Foerster, Justin Gilmer, Jan Chorowski, Jascha Sohl-Dickstein, David Sussillo
- #64 StingyCD: Safely Avoiding Wasteful Updates in Coordinate Descent**
Tyler Johnson, Carlos Guestrin
- #65 Contextual Decision Processes with low Bellman rank are PAC-Learnable**
Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert Schapire
- #66 Tensor Belief Propagation**
Andrew Wrigley, Wee Sun Lee, Nan Ye
- #67 Deep Generative Models for Relational Data with Side Information**
Changwei Hu, Piyush Rai, Lawrence Carin
- #68 Doubly Accelerated Methods for Faster CCA and Generalized Eigendecomposition**
Zeyuan Allen-Zhu, Yuanzhi Li
- #69 Multilevel Clustering via Wasserstein Means**
Nhat Ho, Long Nguyen, Mikhail Yurochkin, Hung Bui, Viet Huynh, Dinh Phung
- #70 Online and Linear-Time Attention by Enforcing Monotonic Alignments**
Colin Raffel, Thang Luong, Peter Liu, Ron Weiss, Douglas Eck
- #71 Stochastic modified equations and adaptive stochastic gradient algorithms**
Qianxiao Li, Cheng Tai, Weinan E
- #72 A Simple Multi-Class Boosting Framework with Theoretical Guarantees and Empirical Proficiency**
Ron Appel, Pietro Perona
- #73 Faster Greedy MAP Inference for Determinantal Point Processes**
Insu Han, Prabhanjan Kambadur, Kyoungsoo Park, Jinwoo Shin
- #74 ChoiceRank: Identifying Preferences from Node Traffic in Networks**
Ium Maystre, Matt Grossglauser
- #75 On the Iteration Complexity of Support Recovery via Hard Thresholding Pursuit**
Jie Shen, Ping Li



- #76 Uniform Deviation Bounds for k-Means Clustering**
Olivier Bachem, Mario Lucic, Hamed Hassani, Andreas Krause
- #77 Sequence Tutor: Conservative fine-tuning of sequence generation models with KL-control**
Natasha Jaques, Shixiang Gu, Dzmitry Bahdanau, Jose Hernandez-Lobato, Richard E Turner, Douglas Eck
- #78 Dissipativity Theory for Nesterov's Accelerated Method**
Bin Hu, Laurent Lessard
- #79 Gradient Boosted Decision Trees for High Dimensional Sparse Output**
Si Si, Huan Zhang, Sathiya Keerthi, Dhruv Mahajan, Inderjit Dhillon, Cho-Jui Hsieh
- #80 Zonotope hit-and-run for efficient sampling from projection DPPs**
Guillaume Gautier, Rémi Bardenet, Michal Valko
- #81 Statistical Inference for Incomplete Ranking Data: The Case of Rank-Dependent Coarsening**
Mohsen Ahmadi Fahandar, Eyke Hüllermeier, Ines Couso
- #82 Dual Iterative Hard Thresholding: From Non-convex Sparse Minimization to Non-smooth Concave Maximization**
Bo Liu, Xiaotong Yuan, Lezi Wang, Qingshan Liu, Dimitris Metaxas
- #83 Uniform Convergence Rates for Kernel Density Estimation**
Heinrich Jiang
- #84 Deep Voice: Real-time Neural Text-to-Speech**
agibiansky Gibiansky, Mike Chrzanowski, Mohammad Shoeybi, Shubho Sengupta, Gregory Diamos, Sercan Arik, Jonathan Raiman, John Miller, Xian Li, Yongguo Kang, Adam Coates, Andrew Ng
- #85 An Analytical Formula of Population Gradient for two-layered ReLU network and its Applications in Convergence and Critical Point Analysis**
Yuandong Tian
- #86 Globally Induced Forest: A Pruning Compression Scheme**
Jean-Michel Begon, Arnaud Joly, Pierre Geurts
- #87 A Divergence Bound for Hybrids of MCMC and Variational Inference and an Application to Langevin Dynamics and SGVI**
Justin Domke
- #88 Just Sort It! A Simple and Effective Approach to Active Preference Learning**
Ium Maystre, Matt Grossglauser
- #89 On The Projection Operator to A Three-view Cardinality Constrained Set**
Haichuan Yang, Shupeng Gui, Chuyang Ke, Daniel Stefankovic, Ryohei Fujimaki, Ji Liu
- #90 Density Level Set Estimation on Manifolds with DBSCAN**
Heinrich Jiang
- #91 DeepBach: a Steerable Model for Bach Chorales Generation**
Gaëtan HADJERES, François Pachet, Frank Nielsen
- #92 Forward and Reverse Gradient-Based Hyperparameter Optimization**
Luca Franceschi, Michele Donini, Paolo Frasconi, Massimiliano Pontil
- #93 Forest-type Regression with General Losses and Robust Forest**
Hanbo Li, Andy Martin
- #94 On the Sampling Problem for Kernel Quadrature**
Francois-Xavier Briol, Chris J Oates, Jon Cockayne, Wilson Ye Chen, Mark Girolami
- #95 Maximum Selection and Ranking under Noisy Comparisons**
Moein Falahatgar, Alon Orlitsky, Venkatadheeraj Pichapati, Ananda Suresh
- #96 Sparse + Group-Sparse Dirty Models: Statistical Guarantees without Unreasonable Conditions and a Case for Non-Convexity**
Eunho Yang, Aurelie Lozano
- #97 Algorithmic Stability and Hypothesis Complexity**
Tongliang Liu, Gábor Lugosi, Gergely Neu, Dacheng Tao
- #98 Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders**
Cinjon Resnick, Adam Roberts, Jesse Engel, Douglas Eck, Sander Dieleman, Karen Simonyan, Mohammad Norouzi
- #99 Adaptive Sampling Probabilities for Non-Smooth Optimization**
Hongseok Namkoong, Aman Sinha, Steve Yadlowsky, John Duchi
- #100 Confident Multiple Choice Learning**
Kimin Lee, Changho Hwang, Kyoungsoo Park, Jinwoo Shin
- #101 Measuring Sample Quality with Kernels**
Jackson Gorham, Lester Mackey
- #102 Active Learning for Top- k Rank Aggregation from Noisy Comparisons**
Soheil Mohajer, Changho Suh, Adel Elmahdy
- #103 Compressed Sensing using Generative Models**
Ashish Bora, Ajil Jalal, Eric Price, Alex Dimakis
- #104 Consistency Analysis for Binary Classification Revisited**
Krzysztof Dembczynski, Wojciech Kotlowski, Sanmi Koyejo, Nagarajan Natarajan
- #105 A Closer Look at Memorization in Deep Networks**
David Krueger, Yoshua Bengio, Stanislaw Jastrzebski, Maxinder S. Kanwal, Nicolas Ballas, Asja Fischer, Emmanuel Bengio, Devansh Arpit, Tegan Maharaj, Aaron Courville, Simon Lacoste-Julien
- #106 Learning to Generate Long-term Future via Hierarchical Prediction**
Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, Honglak Lee
- #107 Sub-sampled Cubic Regularization for Non-convex Optimization**
Jonas Kohler, Aurelien Lucchi
- #108 Regret Minimization in Behaviorally-Constrained Zero-Sum Games**
Gabriele Farina, Christian Kroer, Tuomas Sandholm

- #109 Variational Boosting: Iteratively Refining Posterior Approximations**
Andrew Miller, Nick J Foti, Ryan Adams
- #110 Learning to Align the Source Code to the Compiled Object Code**
Ariel Green, Lior Wolf
- #111 Scaling Up Sparse Support Vector Machines by Simultaneous Feature and Sample Reduction**
Weizhong Zhang, Bin Hong, Wei Liu, Jieping Ye, Deng Cai, Xiaofei He, Jie Wang
- #112 Distributed Mean Estimation with Limited Communication**
Ananda Suresh, Felix Yu, Sanjiv Kumar, Brendan McMahan
- #113 Cognitive Psychology for Deep Neural Networks: A Shape Bias Case Study**
Sam Ritter, David GT Barrett, Adam Santoro, Matthew Botvinick
- #114 Sequence to Better Sequence: Continuous Revision of Combinatorial Structures**
Jonas Mueller, David Gifford, Tommi Jaakkola
- #115 Natasha: Faster Non-Convex Stochastic Optimization Via Strongly Non-Convex Parameter**
Zeyuan Allen-Zhu
- #116 Reduced Space and Faster Convergence in Imperfect-Information Games via Pruning**
Noam Brown, Tuomas Sandholm
- #117 Lost Relatives of the Gumbel Trick**
Matej Balog, Nilesh Tripuraneni, Zoubin Ghahramani, Adrian Weller
- #118 RobustFill: Neural Program Learning under Noisy I/O**
Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdelrahman Mohammad, Pushmeet Kohli
- #119 Efficient Distributed Learning with Sparsity**
Jialei Wang, Mladen Kolar, Nati Srebro, Tong Zhang
- #120 Nonparanormal Information Estimation**
Shashank Singh, Barnabás Póczos
- #121 Visualizing and Understanding Multilayer Perceptron Models: A Case Study in Speech Processing**
Tasha Nagamine, Nima Mesgarani
- #122 Tensor-Train Recurrent Neural Networks for Video Classification**
Yinchong Yang, Denis Krompass, Volker Tresp
- #123 “Convex Until Proven Guilty”: Dimension-Free Acceleration of Gradient Descent on Non-Convex Functions**
Yair Carmon, John Duchi, Oliver Hinder, Aaron Sidford
- #124 Strongly-Typed Agents are Guaranteed to Interact Safely**
David Balduzzi
- #125 Learning to Aggregate Ordinal Labels by Maximizing Separating Width**
Guangyong Chen, Shengyu Zhang, Di Lin, HUI Huang, Pheng Heng
- #126 Programming with a Differentiable Forth Interpreter**
Matko Bošnjak, Tim Rocktäschel, Jason Naradowsky, Sebastian Riedel
- #127 Innovation Pursuit: A New Approach to the Subspace Clustering Problem**
Mostafa Rahmani, George Atia
- #128 A Unified Maximum Likelihood Approach for Estimating Symmetric Properties of Discrete Distributions**
Jayadev Acharya, Hirakendu Das, Alon Orlitsky, Ananda Suresh
- #129 Axiomatic Attribution for Deep Networks**
Mukund Sundararajan, Ankur Taly, Qiqi Yan
- #130 Sequence Modeling via Segmentations**
Chong Wang, Yining Wang, Po-Sen Huang, Abdelrahman Mohammad, Dengyong Zhou, Li Deng
- #131 Convergence Analysis of Proximal Gradient with Momentum for Nonconvex Optimization**
Qunwei Li, Yi Zhou, Yingbin Liang, Pramod K Varshney
- #132 Coordinated Multi-Agent Imitation Learning**
Hoang Le, Yisong Yue, Peter Carr, Patrick Lucey
- #133 Uncorrelation and Evenness: a New Diversity-Promoting Regularizer**
Pengtao Xie, Aarti Singh, Eric Xing
- #134 Differentiable Programs with Neural Libraries**
Alex Gaunt, Marc Brockschmidt, Nate Kushman, Daniel Tarlow
- #135 Selective Inference for Sparse High-Order Interaction Models**
Shinya Suzumura, Kazuya Nakagawa, Yuta Umezu, Koji Tsuda, Ichiro Takeuchi
- #136 Gradient Coding: Avoiding Stragglers in Distributed Learning**
Rashish Tandon, Qi Lei, Alex Dimakis, NIKOS KARAMPATZIAKIS
- #137 On Calibration of Modern Neural Networks**
Chuan Guo, Geoff Pleiss, Yu Sun, Kilian Weinberger
- #138 Latent LSTM Allocation: Joint clustering and non-linear dynamic modeling of sequence data**
Manzil Zaheer, Amr Ahmed, Alex Smola
- #139 How to Escape Saddle Points Efficiently**
Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M. Kakade, Michael Jordan
- #140 Deep Decentralized Multi-task Multi-Agent Reinforcement Learning under Partial Observability**
Shayegan Omidshafiei, Jason Pazis, Chris Amato, Jonathan How, John L Vian
- #141 Learning Latent Space Models with Angular Constraints**
Pengtao Xie, Yuntian Deng, Yi Zhou, Abhimanu Kumar, Yaoliang Yu, James Zou, Eric Xing
- #142 Developing Bug-Free Machine Learning Systems With Formal Mathematics**
Daniel Selsam, Percy Liang, David L Dill
- #143 Dictionary Learning Based on Sparse Distribution Tomography**
Pedram Pad, Farnood Salehi, Elisa Celis, Patrick Thiran, Michael Unser
- #144 Learning Discrete Representations via Information Maximizing Self-Augmented Training**
Weihua Hu, Takeru Miyato, Seiya Tokui, Eiichi Matsumoto, Masashi Sugiyama

Wednesday

AUGUST 9TH | SESSIONS



TIME	DESCRIPTION	LOCATION	TIME	DESCRIPTION	LOCATION
9 am	Invited Talk: Raia Hadsell Towards Reinforcement Learning in the Real World	L2, Darling Harbour	3 pm	Break	
10 am	Break		3:30 pm	SESSION THREE	
10:30 am	SESSION ONE			Applications	L2, Darling Harbour
	Deep learning 8: hardware	L2, Darling Harbour		Language 3	L2, Parkside 1
	Language 1	L2, Parkside 1		Combinatorial optimization 2	L2, Parkside 2
	Distributed optimization	L2, Parkside 2		Deep reinforcement learning 2	L4, C4.5
	Continuous control	L4, C4.5		Gaussian processes	L4, C4.9 & C4.10
	Monte Carlo methods 1	L4, C4.9 & C4.10		Causal Inference 2	L4, C4.1
	Bayesian Optimization	L4, C4.1		Large scale learning	L4, C4.4
	High dimensional estimation	L4, C4.4		Kernel methods	L4, C4.8
	Privacy and security 1	L4, C4.8		Structured prediction	L4, C4.6 & C4.7
	Healthcare	L4, C4.6 & C4.7	5:15 pm	Invited Talk: Bernhard Schölkopf	
12 pm	LUNCH (On Your Own)			Causal Learning	L2, Darling Harbour
1:30 pm	SESSION TWO		6:15 pm	Poster Session	L2, Gallery
	Deep learning 9: probabilistic	L2, Darling Harbour			
	Language 2	L2, Parkside 1			
	Combinatorial optimization 1	L2, Parkside 2			
	Deep reinforcement learning 1	L4, C4.5			
	Monte Carlo methods 2	L4, C4.9 & C4.10			
	Causal Inference 1	L4, C4.1			
	Robust Estimation	L4, C4.4			
	Privacy and security 2	L4, C4.8			
	Time series	L4, C4.6 & C4.7			



SESSION ONE - 10:30 AM - 12:00 PM

Deep learning 8: hardware

Location: Level 2, Darling Harbour

- **Device Placement Optimization with Reinforcement Learning**
Azalia Mirhoseini, Hieu Pham, Quoc Le, benoit steiner, Mohammad Norouzi, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Samy Bengio, Jeff Dean
- **Deep Tensor Convolution on Multicores**
David Budden, Alexander Matveev, Shibani Santurkar, Shraman Ray Chaudhuri, Nir Shavit
- **MEC: Memory-efficient Convolution for Deep Neural Network**
Minsik Cho, Daniel Brand
- **Beyond Filters: Compact Feature Map for Portable Deep Model**
Yunhe Wang, Chang Xu, Chao Xu, Dacheng Tao
- **Efficient softmax approximation for GPUs**
Edouard Grave, Armand Joulin, Moustapha Cisse, David Grangier, Herve Jegou

Language 1

Location: Level 2, Parkside 1

- **Dynamic Word Embeddings**
Robert Bamler, Stephan Mandt
- **Gram-CTC: Automatic Unit Selection and Target Decomposition for Sequence Labelling**
Hairong Liu, Zhenyao Zhu, Xiangang Li, Sanjeev Satheesh
- **Coupling Distributed and Symbolic Execution for Natural Language Queries**
Lili Mou, Zhengdong Lu, Hang Li, Zhi Jin
- **Image-to-Markup Generation with Coarse-to-Fine Attention**
Yuntian Deng, Anssi Kanervisto, Jeffrey Ling, Alexander Rush
- **Multichannel End-to-end Speech Recognition**
Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, John Hershey

Distributed optimization

Location: Level 2, Parkside 2

- **Asynchronous Stochastic Gradient Descent with Delay Compensation**
Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhiming Ma, Tie-Yan Liu
- **Adaptive Consensus ADMM for Distributed Optimization**
Zheng Xu, Gavin Taylor, Hao Li, Mario Figueiredo, Xiaoming Yuan, Tom Goldstein
- **Optimal Algorithms for Smooth and Strongly Convex Distributed Optimization in Networks**
Kevin Scaman, Francis Bach, Sebastien Bubeck, Yin Tat Lee, Laurent Massoulié
- **Projection-free Distributed Online Learning in Networks**
Wenpeng Zhang, Peilin Zhao, wenwu zhu, Steven Hoi, Tong Zhang

Continuous control

Location: Level 4, C4.5

- **Improving Stochastic Policy Gradients in Continuous Control with Deep Reinforcement Learning using the Beta Distribution**
Po-Wei Chou, Daniel Maturana, Sebastian Scherer
- **Combining Model-Based and Model-Free Updates for Trajectory-Centric Reinforcement Learning**
Yevgen Chebotar, Karol Hausman, Marvin Zhang, Gaurav Sukhatme, Stefan Schaal, Sergey Levine
- **Prediction under Uncertainty in Sparse Spectrum Gaussian Processes with Applications to Filtering and Control**
Yunpeng Pan, Xinyan Yan, Evangelos Theodorou, Byron Boots
- **Learning Stable Stochastic Nonlinear Dynamical Systems**
Jonas Umlauft, Sandra Hirche
- **Local Bayesian Optimization of Motor Skills**
Riad Akrou, Dmitry Sorokin, Jan Peters, Gerhard Neumann



SESSION ONE - 10:30 AM - 12:00 PM

Monte Carlo methods 1

Location: Level 4, C4.9& C4.10

- **Fractional Langevin Monte Carlo: Exploring Levy Driven Stochastic Differential Equations for MCMC**
Umut Simsekli
- **Stochastic Bouncy Particle Sampler**
Ari Pakman, Dar Gilboa, David Carlson, Liam Paninski
- **Canopy --- Fast Sampling with Cover Trees**
Manzil Zaheer, Satwik Kottur, Amr Ahmed, Jose Moura, Alex Smola
- **A Simulated Annealing Based Inexact Oracle for Wasserstein Loss Minimization**
Jianbo Ye, James Wang, Jia Li
- **Improving Gibbs Sampler Scan Quality with DoGS**
Ioannis Mitliagkas, Lester Mackey

Bayesian Optimization

Location: Level 4, C4.1

- **Preferential Bayesian Optimization**
Javier González, Zhenwen Dai, Andreas Damianou, Neil Lawrence
- **Max-value Entropy Search for Efficient Bayesian Optimization**
Zi Wang, Stefanie Jegelka
- **Bayesian Optimization with Tree-structured Dependencies**
Rodolphe Jenatton, Cedric Archambeau, Javier González, Matthias Seeger
- **Multi-fidelity Bayesian Optimisation with Continuous Approximations**
kirthevasan kandasamy, Gautam Dasarathy, Barnabás Póczos, Jeff Schneider
- **Parallel and Distributed Thompson Sampling for Large-scale Accelerated Exploration of Chemical Space**
Jose Hernandez-Lobato, James Requeima, Edward Pyzer-Knapp, alan Aspuru-Guzik

High dimensional estimation

Location: Level 4, C4.4

- **Being Robust (in High Dimensions) Can Be Practical**
Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, Alistair Stewart
- **Multilabel Classification with Group Testing and Codes**
Shashanka Ubaru, Arya Mazumdar
- **High-Dimensional Structured Quantile Regression**
Vidyashankar Sivakumar, Arindam Banerjee

- **High-dimensional Non-Gaussian Single Index Models via Thresholded Score Function Estimation**
Zhuoran Yang, Krishnakumar Balasubramanian, Han Liu
- **Robust Structured Estimation with Single-Index Models**
Sheng Chen, Arindam Banerjee

Privacy and security 1

Location: Level 4, C4.8

- **Differentially Private Ordinary Least Squares**
Or Sheffet
- **Priv'IT: Private and Sample Efficient Identity Testing**
Bryan Cai, Constantinos Daskalakis, Gautam Kamath
- **Differentially Private Submodular Maximization: Data Summarization in Disguise**
Marko Mitrovic, Mark Bun, Andreas Krause, Amin Karbasi
- **Differentially Private Learning of Graphical Models using CGMs**
Garrett Bernstein, Ryan McKenna, Tao Sun, Daniel Sheldon, Michael Hay, Gerome Miklau
- **Minimizing Trust Leaks for Robust Sybil Detection**
János Höner, Shinichi Nakajima, Alexander Bauer, Klaus-robot Mueller, Nico Görnitz

Healthcare

Location: Level 4, C4.6 & C4.7

- **When can Multi-Site Datasets be Pooled for Regression? Hypothesis Tests, ℓ_2 -consistency and Neuroscience Applications**
Hao Zhou, Yilin Zhang, Vamsi Ithapu, Sterling Johnson, Grace Wahba, Vikas Singh
- **Learning from Clinical Judgments: Semi-Markov-Modulated Marked Hawkes Processes for Risk Prognosis**
Ahmed M. Alaa Ibrahim, Scott B Hu, Mihaela van der Schaar
- **Learning to Detect Sepsis with a Multitask Gaussian Process RNN Classifier**
Joseph Futoma, Sanjay Hariharan, Katherine Heller
- **iSurvive: An Interpretable, Event-time Prediction Model for mHealth**
Walter Dempsey, Alexander Moreno, Jim Rehg, Susan Murphy, Chris Scott, Michael Dennis, David Gustafson
- **Learning Sleep Stages from Radio Signals: A Conditional Adversarial Architecture**
Mingmin Zhao, Shichao Yue, Dina Katabi, Tommi Jaakkola, Matt Bianchi



SESSION TWO - 1:30 - 3:00 PM

Deep learning 9: probabilistic

Location: Level 2, Darling Harbour

- **Dropout Inference in Bayesian Neural Networks with Alpha-divergences**
Yingzhen Li, Yarin Gal
- **Multiplicative Normalizing Flows for Variational Bayesian Neural Networks**
Christos Louizos, Max Welling
- **Variational Dropout Sparsifies Deep Neural Networks**
Dmitry Molchanov, Arsenii Ashukha, Dmitry Vetrov
- **Unimodal Probability Distributions for Deep Ordinal Classification**
Christopher Beckham, Christopher Pal

Language 2

Location: Level 2, Parkside 1

- **Latent Intention Dialogue Models**
Tsung-Hsien Wen, Yishu Miao, Philip Blunsom, Steve Young J Young
- **Discovering Discrete Latent Topics with Neural Variational Inference**
Yishu Miao, Edward Grefenstette, Philip Blunsom
- **Toward Controlled Generation of Text**
Zhiting Hu, Zichao Yang, Xiaodan Liang, Russ Salakhutdinov, Eric Xing
- **Learning Continuous Semantic Representations of Symbolic Expressions**
Miltos Allamanis, pankajan Chanthirasegaran, Pushmeet Kohli, Charles Sutton
- **Adversarial Feature Matching for Text Generation**
Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, Lawrence Carin

Combinatorial optimization 1

Location: Level 2, Parkside 2

- **Robust Guarantees of Stochastic Greedy Algorithms**
Yaron Singer, Avinandan Hassidim
- **Guarantees for Greedy Maximization of Non-submodular Functions with Applications**
Andrew An Bian, Joachim Buhmann, Andreas Krause, Sebastian Tschiatschek
- **Robust Submodular Maximization: A Non-Uniform Partitioning Approach**
Ilija Bogunovic, Boba Mitrovic, Jonathan Scarlett, Volkan Cevher

- **Probabilistic Submodular Maximization in Sub-Linear Time**
Serban A Stan, Morteza Zadimoghaddam, Andreas Krause, Amin Karbasi
- **On Approximation Guarantees for Greedy Low Rank Optimization**
RAJIV KHANNA, Ethan Elenberg, Alex Dimakis, Joydeep Ghosh, Sahand Negahban

Deep reinforcement learning 1

Location: Level 4, C4.5

- **Count-Based Exploration with Neural Density Models**
Georg Ostrovski, Marc Bellemare, Aäron van den Oord, Remi Munos
- **Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning**
Junhyuk Oh, Satinder Singh, Honglak Lee, Pushmeet Kohli
- **Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning**
Jakob Foerster, Nantas Nardelli, Greg Farquhar, Triantafyllos Afouras, Phil Torr, Pushmeet Kohli, Shimon Whiteson
- **The Predictron: End-To-End Learning and Planning**
David Silver, Hado van Hasselt, Matteo Hessel, Tom Schaul, Arthur Guez, Tim Harley, Gabriel Dulac-Arnold, David Reichert, Neil Rabinowitz, Andre Barreto, Thomas Degris
- **Averaged-DQN: Variance Reduction and Stabilization for Deep Reinforcement Learning**
Oron Anschel, Nir Baram, Nahum Shimkin

Monte Carlo methods 2

Location: Level 4, C4.9& C4.10

- **Magnetic Hamiltonian Monte Carlo**
Nilesh Tripuraneni, Mark Rowland, Zoubin Ghahramani, Richard E Turner
- **Probabilistic Path Hamiltonian Monte Carlo**
Vu Dinh, Arman Bilge, Cheng Zhang, Erick Matsen
- **Stochastic Gradient Monomial Gamma Sampler**
Yizhe Zhang, Changyou Chen, Zhe Gan, Ricardo Henao, Lawrence Carin
- **Stochastic Gradient MCMC Methods for Hidden Markov Models**
Yi-An Ma, Nick J Foti, Emily Fox
- **Deep Latent Dirichlet Allocation with Topic-Layer-Adaptive Stochastic Gradient Riemannian MCMC**
Yulai Cong, Bo Chen, Hongwei Liu, Mingyuan Zhou



SESSION TWO - 1:30 - 3:00 PM

Causal Inference 1

Location: Level 4, C4.1

- **Uncertainty Assessment and False Discovery Rate Control in High-Dimensional Granger Causal Inference**
Aditya Chaudhry, Pan Xu, Quanquan Gu
- **Uncovering Causality from Multivariate Hawkes Integrated Cumulants**
Massil Achab, Emmanuel Bacry, Stéphane Gaïffas, Jacopo Mastromatteo, Jean-François Muzy
- **Cost-Optimal Learning of Causal Graphs**
Murat Kocaoglu, Alex Dimakis, Sriram Vishwanath
- **Identification and Model Testing in Linear Structural Equation Models using Auxiliary Variables**
Bryant Chen, Daniel Kumor, Elias Bareinboim
- **Estimating individual treatment effect: generalization bounds and algorithms**
Uri Shalit, Fredrik D Johansson, David Sontag

Robust Estimation

Location: Level 4, C4.4

- **Toward Efficient and Accurate Covariance Matrix Estimation on Compressed Data**
Xixian Chen, Michael Lyu, Irwin King
- **Robust Gaussian Graphical Model Estimation with Arbitrary Corruption**
Lingxiao Wang, Quanquan Gu
- **Algebraic Variety Models for High-Rank Matrix Completion**
Greg Ongie, Laura Balzano, Rebecca Willett, Robert Nowak
- **High-Dimensional Variance-Reduced Stochastic Gradient Expectation-Maximization Algorithm**
Rongda Zhu, Lingxiao Wang, Chengxiang Zhai, Quanquan Gu

Privacy and security 2

Location: Level 4, C4.8

- **The Price of Differential Privacy For Online Learning**
Naman Agarwal, Karan Singh
- **Pain-Free Random Differential Privacy with Sensitivity Sampling**
Ben Rubinstein, Francesco Aldà
- **Differentially Private Clustering in High-Dimensional Euclidean Spaces**
Nina Balcan, Travis Dick, Yingyu Liang, Wenlong Mou, Hongyang Zhang
- **Differentially Private Chi-squared Test by Unit Circle Mechanism**
Kazuya Kakizaki, Kazuto Fukuchi, Jun Sakuma
- **Collect at Once, Use Effectively: Making Non-interactive Locally Private Learning Possible**
Kai Zheng, Wenlong Mou, Liwei Wang

Time series

Location: Level 4, C4.6 & C4.7

- **Bidirectional learning for time-series models with hidden units**
Takayuki Osogami, Hiroshi Kajino, Taro Sekiyama
- **Learning Hawkes Processes from Short Doubly-Censored Event Sequences**
Hongteng Xu, Dixin Luo, Hongyuan Zha
- **Coherent probabilistic forecasts for hierarchical time series**
Souhaib Ben Taieb, James Taylor, Rob Hyndman
- **Soft-DTW: a Differentiable Loss Function for Time-Series**
Marco Cuturi, Mathieu Blondel
- **Variational Policy for Guiding Point Processes**
Yichen Wang, Grady Williams, Evangelos Theodorou, Le Song



SESSION THREE - 3:30 - 5:00 PM

Applications

Location: Level 2, Darling Harbour

- **Dance Dance Convolution**
Chris Donahue, Zachary Lipton, Julian McAuley
- **World of Bits: An Open-Domain Platform for Web-Based Agents**
Tim Shi, Andrej Karpathy, Jim Fan Fan, Jonathan Hernandez, Percy Liang
- **Real-Time Adaptive Image Compression**
Oren Rippel, Lubomir Bourdev
- **Neural Message Passing for Quantum Chemistry**
Justin Gilmer, Samuel Schoenholz, Patrick F Riley, Oriol Vinyals, George Dahl
- **Accelerating Eulerian Fluid Simulation With Convolutional Networks**
Jonathan Tompson, Kristofer D Schlachter, Pablo Sprechmann, Ken Perlin

Language 3

Location: Level 2, Parkside 1

- **Language Modeling with Gated Convolutional Networks**
Yann Dauphin, Angela Fan, Michael Auli, David Grangier
- **Convolutional Sequence to Sequence Learning**
Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, Yann Dauphin
- **Improved Variational Autoencoders for Text Modeling using Dilated Convolutions**
Zichao Yang, Zhiting Hu, Russ Salakhutdinov, Taylor Berg-Kirkpatrick
- **Grammar Variational Autoencoder**
Matt J. Kusner, Brooks Paige, Jose Hernandez-Lobato

Combinatorial optimization 2

Location: Level 2, Parkside 2

- **Deletion-Robust Submodular Maximization: Data Summarization with “the Right to be Forgotten”**
Baharan Mirzasoleiman, Amin Karbasi, Andreas Krause
- **Analysis and Optimization of Graph Decompositions by Lifted Multicuts**
Andrea Hornakova, Jan-Hendrik Lange, Bjoern Andres

- **Near-Optimal Design of Experiments via Regret Minimization**
Zeyuan Allen-Zhu, Yuanzhi Li, Aarti Singh, Yining Wang
- **Robust Budget Allocation via Continuous Submodular Functions**
Matthew J Staib, Stefanie Jegelka
- **Rule-Enhanced Penalized Regression by Column Generation using Rectangular Maximum Agreement**
Jonathan Eckstein, Noam Goldberg, Ai Kagawa

Deep reinforcement learning 2

Location: Level 4, C4.5

- **FeUdal Networks for Hierarchical Reinforcement Learning**
Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, Koray Kavukcuoglu
- **Deciding How to Decide: Dynamic Routing in Artificial Neural Networks**
Mason McGill, Pietro Perona
- **Neural Episodic Control**
Alexander Pritzel, Benigno Uria, Raam Sriram, Adrià Puigdomenech Badia, Oriol Vinyals, Demis Hassabis, Daan Wierstra, Charles Blundell
- **Neural Optimizer Search using Reinforcement Learning**
Irwan Bello, Barret Zoph, Vijay Vasudevan, Quoc Le

Gaussian processes

Location: Level 4, C4.9& C4.10

- **Distributed Batch Gaussian Process Optimization**
Erik Daxberger, Bryan Kian Hsiang Low
- **Scalable Multi-Class Gaussian Process Classification using Expectation Propagation**
Carlos Villacampa-Calvo, Daniel Hernandez-Lobato
- **Random Feature Expansions for Deep Gaussian Processes**
Kurt Cutajar, Edwin Bonilla, Pietro Michiardi, Maurizio Filippone
- **Asynchronous Distributed Variational Gaussian Processes for Regression**
Hao Peng, Shandian Zhe, Cosmo Zhang, Alan Qi
- **High Dimensional Bayesian Optimization with Elastic Gaussian Process**
Santu Rana, Cheng Li, Sunil Gupta, Vu Nguyen, Svetha Venkatesh



SESSION THREE - 3:30 - 5:00 PM

Causal Inference 2

Location: Level 4, C4.1

- **Recursive Partitioning for Personalization using Observational Data**
Nathan Kallus
- **Identifying Best Interventions through Online Importance Sampling**
Rajat Sen, Karthikeyan Shanmugam, Alex Dimakis, Sanjay Shakkottai
- **Deep IV: A Flexible Approach for Counterfactual Prediction**
Jason Hartford, Greg Lewis, Kevin Leyton-Brown, Matt Taddy
- **Counterfactual Data-Fusion for Online Reinforcement Learners**
Andrew Forney, Judea Pearl, Elias Bareinboim

Large scale learning

Location: Level 4, C4.4

- **Optimal Densification for Fast and Accurate Minwise Hashing**
Anshumali Shrivastava
- **Stochastic Generative Hashing**
Bo Dai, Ruiqi Guo, Sanjiv Kumar, Niao He, Le Song
- **ZipML: Training Linear Models with End-to-End Low Precision, and a Little Bit of Deep Learning**
Hantian Zhang, Jerry Li, Kaan Kara, Dan Alistarh, Ji Liu, Ce Zhang
- **Large-Scale Evolution of Image Classifiers**
Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc Le, Alex Kurakin

Kernel methods

Location: Level 4, C4.8

- **An Adaptive Test of Independence with Analytic Kernel Embeddings**
Wittawat Jitkrittum, Zoltan Szabo, Arthur Gretton
- **Sliced Wasserstein Kernel for Persistence Diagrams**
Mathieu Carrière, Marco Cuturi, Steve Oudot
- **Adapting Kernel Representations Online Using Submodular Maximization**
Matthew Schlegel, Yangchen Pan, Jiecao (Jack) Chen, Martha White
- **Spherical Structured Feature Maps for Kernel Approximation**
Yueming LYU
- **Nyström Method with Kernel K-means++ Samples as Landmarks**
Dino Ogliz, Thomas Gaertner

Structured prediction

Location: Level 4, C4.6 & C4.7

- **Deep Value Networks Learn to Evaluate and Iteratively Refine Structured Outputs**
Gygli Gygli, Mohammad Norouzi, Anelia Angelova
- **Deeply AggreVaTeD: Differentiable Imitation Learning for Sequential Prediction**
Wen Sun, Arun Venkatraman, Geoff Gordon, Byron Boots, Drew Bagnell
- **End-to-End Learning for Structured Prediction Energy Networks**
David Belanger, Bishan Yang, Andrew McCallum
- **A Unified View of Multi-Label Performance Measures**
Xi-Zhu Wu, Zhi-Hua Zhou
- **Scalable Generative Models for Multi-label Learning with Missing Labels**
Vikas Jain, Nirbhay Modhe, Piyush Rai



- #1 Sketched Ridge Regression: Optimization Perspective, Statistical Perspective, and Model Averaging**
Shusen Wang, Alex Gittens, Michael Mahoney
- #2 Estimating the unseen from multiple populations**
Aditi Raghunathan, Greg Valiant, James Zou
- #3 Meritocratic Fairness for Cross-Population Selection**
Michael Kearns, Aaron Roth, Steven Wu
- #4 Neural networks and rational functions**
Matus Telgarsky
- #5 Input Convex Neural Networks**
Brandon Amos, Lei Xu, Zico Kolter
- #6 Co-clustering through Optimal Transport**
Charlotte Laclau, Ievgen Redko, Basarab Matej, Younès Bennani, Vincent Brault
- #7 OptNet: Differentiable Optimization as a Layer in Neural Networks**
Brandon Amos, Zico Kolter
- #8 Multiple Clustering Views from Multiple Uncertain Experts**
Yale Chang, Junxiang Chen, Michael Cho, Peter Castaldi, Edwin Silverman, Jennifer G Dy
- #9 Parseval Networks: Improving Robustness to Adversarial Examples**
Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, Nicolas Usunier
- #10 Clustering by Sum of Norms: Stochastic Incremental Algorithm, Convergence and Cluster Recovery**
Ashkan Panahi, Devdatt Dubhashi, Fredrik D Johansson, Chiranjib Bhattacharya
- #11 Regularising Non-linear Models Using Feature Side-information**
Amina Mollaysa, Pablo Strasser, Alexandros Kalousis
- #12 Clustering High Dimensional Dynamic Data Streams**
Lin Yang, Harry Lang, Christian Sohler, Vladimir Braverman, Gereon Frahling
- #13 Fast k-Nearest Neighbour Search via Prioritized DCI**
Ke Li, Jitendra Malik
- #14 Deep Spectral Clustering Learning**
Marc Law, Raquel Urtasun, Zemel Rich
- #15 Joint Dimensionality Reduction and Metric Learning: A Geometric Take**
Mehrtash Harandi, Mathieu Salzmann, Richard I Hartley
- #16 ProtoNN: Compressed and Accurate kNN for Resource-scarce Devices**
Chirag Gupta, ARUN SUGGALA, Ankit Goyal, Saurabh Goyal, Ashish Kumar, Bhargavi Paranjape, Harsha Vardhan Simhadri, Raghavendra Udupa, Manik Varma, Prateek Jain
- #17 Device Placement Optimization with Reinforcement Learning**
Azalia Mirhoseini, Hieu Pham, Quoc Le, benoit steiner, Mohammad Norouzi, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Samy Bengio, Jeff Dean
- #18 Dynamic Word Embeddings**
Robert Bamler, Stephan Mandt
- #19 Asynchronous Stochastic Gradient Descent with Delay Compensation**
Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhiming Ma, Tie-Yan Liu
- #20 Improving Stochastic Policy Gradients in Continuous Control with Deep Reinforcement Learning using the Beta Distribution**
Po-Wei Chou, Daniel Maturana, Sebastian Scherer
- #21 Fractional Langevin Monte Carlo: Exploring Levy Driven Stochastic Differential Equations for MCMC**
Umut Simsekli
- #22 Preferential Bayesian Optimization**
Javier González, Zhenwen Dai, Andreas Damianou, Neil Lawrence
- #23 Being Robust (in High Dimensions) Can Be Practical**
Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, Alistair Stewart
- #24 Differentially Private Ordinary Least Squares**
Or Sheffet
- #25 When can Multi-Site Datasets be Pooled for Regression? Hypothesis Tests, ℓ_2 -consistency and Neuroscience Applications**
Hao Zhou, Yilin Zhang, Vamsi Ithapu, Sterling Johnson, Grace Wahba, Vikas Singh
- #26 Deep Tensor Convolution on Multicores**
David Budden, Alexander Matveev, Shibani Santurkar, Shraman Ray Chaudhuri, Nir Shavit
- #27 Gram-CTC: Automatic Unit Selection and Target Decomposition for Sequence Labelling**
Hairong Liu, Zhenyao Zhu, Xiangang Li, Sanjeev Satheesh
- #28 Adaptive Consensus ADMM for Distributed Optimization**
Zheng Xu, Gavin Taylor, Hao Li, Mario Figueiredo, Xiaoming Yuan, Tom Goldstein
- #29 Combining Model-Based and Model-Free Updates for Trajectory-Centric Reinforcement Learning**
Yevgen Chebotar, Karol Hausman, Marvin Zhang, Gaurav Sukhatme, Stefan Schaal, Sergey Levine
- #30 Stochastic Bouncy Particle Sampler**
Ari Pakman, Dar Gilboa, David Carlson, Liam Paninski
- #31 Max-value Entropy Search for Efficient Bayesian Optimization**
Zi Wang, Stefanie Jegelka
- #32 Multilabel Classification with Group Testing and Codes**
Shashanka Ubaru, Arya Mazumdar
- #33 Priv'IT: Private and Sample Efficient Identity Testing**
Bryan Cai, Constantinos Daskalakis, Gautam Kamath
- #34 Learning from Clinical Judgments: Semi-Markov-Modulated Marked Hawkes Processes for Risk Prognosis**
Ahmed M. Alaa Ibrahim, Scott B Hu, Mihaela van der Schaar
- #35 MEC: Memory-efficient Convolution for Deep Neural Network**
Minsik Cho, Daniel Brand



- #36 Coupling Distributed and Symbolic Execution for Natural Language Queries**
Lili Mou, Zhengdong Lu, Hang Li, Zhi Jin
- #37 Optimal Algorithms for Smooth and Strongly Convex Distributed Optimization in Networks**
Kevin Scaman, Francis Bach, Sebastien Bubeck, Yin Tat Lee, Laurent Massoulié
- #38 Prediction under Uncertainty in Sparse Spectrum Gaussian Processes with Applications to Filtering and Control**
Yunpeng Pan, Xinyan Yan, Evangelos Theodorou, Byron Boots
- #39 Canopy --- Fast Sampling with Cover Trees**
Manzil Zaheer, Satwik Kottur, Amr Ahmed, Jose Moura, Alex Smola
- #40 Bayesian Optimization with Tree-structured Dependencies**
Rodolphe Jenatton, Cedric Archambeau, Javier González, Matthias Seeger
- #41 High-Dimensional Structured Quantile Regression**
Vidyashankar Sivakumar, Arindam Banerjee
- #42 Differentially Private Submodular Maximization: Data Summarization in Disguise**
Marko Mitrovic, Mark Bun, Andreas Krause, Amin Karbasi
- #43 Learning to Detect Sepsis with a Multitask Gaussian Process RNN Classifier**
Joseph Futoma, Sanjay Hariharan, Katherine Heller
- #44 Beyond Filters: Compact Feature Map for Portable Deep Model**
Yunhe Wang, Chang Xu, Chao Xu, Dacheng Tao
- #45 Image-to-Markup Generation with Coarse-to-Fine Attention**
Yuntian Deng, Anssi Kanervisto, Jeffrey Ling, Alexander Rush
- #46 Projection-free Distributed Online Learning in Networks**
Wenpeng Zhang, Peilin Zhao, wenwu zhu, Steven Hoi, Tong Zhang
- #47 Learning Stable Stochastic Nonlinear Dynamical Systems**
Jonas Umlauft, Sandra Hirche
- #48 A Simulated Annealing Based Inexact Oracle for Wasserstein Loss Minimization**
Jianbo Ye, James Wang, Jia Li
- #49 Multi-fidelity Bayesian Optimisation with Continuous Approximations**
kirthevasan kandasamy, Gautam Dasarathy, Barnabás Póczos, Jeff Schneider
- #50 High-dimensional Non-Gaussian Single Index Models via Thresholded Score Function Estimation**
Zhuoran Yang, Krishnakumar Balasubramanian, Han Liu
- #51 Differentially Private Learning of Graphical Models using CGMs**
Garrett Bernstein, Ryan McKenna, Tao Sun, Daniel Sheldon, Michael Hay, Jerome Miklau
- #52 iSurvive: An Interpretable, Event-time Prediction Model for mHealth**
Walter Dempsey, Alexander Moreno, Jim Rehg, Susan Murphy, Chris Scott, Michael Dennis, David Gustafson
- #53 Efficient softmax approximation for GPUs**
Edouard Grave, Armand Joulin, Moustapha Cisse, David Grangier, Herve Jegou
- #54 Multichannel End-to-end Speech Recognition**
Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, John Hershey
- #55 Local Bayesian Optimization of Motor Skills**
Riad Akrou, Dmitry Sorokin, Jan Peters, Gerhard Neumann
- #56 Improving Gibbs Sampler Scan Quality with DoGS**
Ioannis Mitliagkas, Lester Mackey
- #57 Parallel and Distributed Thompson Sampling for Large-scale Accelerated Exploration of Chemical Space**
Jose Hernandez-Lobato, James Requeima, Edward Pyzer-Knapp, alan Aspuru-Guzik
- #58 Robust Structured Estimation with Single-Index Models**
Sheng Chen, Arindam Banerjee
- #59 Minimizing Trust Leaks for Robust Sybil Detection**
János Höner, Shinichi Nakajima, Alexander Bauer, Klaus-robert Mueller, Nico Görnitz
- #60 Learning Sleep Stages from Radio Signals: A Conditional Adversarial Architecture**
Mingmin Zhao, Shichao Yue, Dina Katabi, Tommi Jaakkola, Matt Bianchi
- #61 Dropout Inference in Bayesian Neural Networks with Alpha-divergences**
Yingzhen Li, Yarin Gal
- #62 Latent Intention Dialogue Models**
Tsung-Hsien Wen, Yishu Miao, Philip Blunsom, Steve Young J Young
- #63 Robust Guarantees of Stochastic Greedy Algorithms**
Yaron Singer, Avinatan Hassidim
- #64 Count-Based Exploration with Neural Density Models**
Georg Ostrovski, Marc Bellemare, Aaron van den Oord, Remi Munos
- #65 Magnetic Hamiltonian Monte Carlo**
Nilesh Tripuraneni, Mark Rowland, Zoubin Ghahramani, Richard E Turner
- #66 Uncertainty Assessment and False Discovery Rate Control in High-Dimensional Granger Causal Inference**
Aditya Chaudhry, Pan Xu, Quanquan Gu
- #67 Toward Efficient and Accurate Covariance Matrix Estimation on Compressed Data**
Xixian Chen, Michael Lyu, Irwin King
- #68 The Price of Differential Privacy For Online Learning**
Naman Agarwal, Karan Singh
- #69 Bidirectional learning for time-series models with hidden units**
Takayuki Osogami, Hiroshi Kajino, Taro Sekiyama
- #70 Multiplicative Normalizing Flows for Variational Bayesian Neural Networks**
Christos Louizos, Max Welling
- #71 Discovering Discrete Latent Topics with Neural Variational Inference**
Yishu Miao, Edward Grefenstette, Philip Blunsom

- #72 Guarantees for Greedy Maximization of Non-submodular Functions with Applications**
Andrew An Bian, Joachim Buhmann, Andreas Krause, Sebastian Tschachtschek
- #73 Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning**
Junhyuk Oh, Satinder Singh, Honglak Lee, Pushmeet Kohli
- #74 Probabilistic Path Hamiltonian Monte Carlo**
Vu Dinh, Arman Bilge, Cheng Zhang, Erick Matsen
- #75 Uncovering Causality from Multivariate Hawkes Integrated Cumulants**
Massil Achab, Emmanuel Bacry, Stéphane Gaïffas, Iacopo Mastromatteo, Jean-François Muzy
- #76 Robust Gaussian Graphical Model Estimation with Arbitrary Corruption**
Lingxiao Wang, Quanquan Gu
- #77 Pain-Free Random Differential Privacy with Sensitivity Sampling**
Ben Rubinfeld, Francesco Aldà
- #78 Learning Hawkes Processes from Short Doubly-Censored Event Sequences**
Hongteng Xu, Dixin Luo, Hongyuan Zha
- #79 Variational Dropout Sparsifies Deep Neural Networks**
Dmitry Molchanov, Arsenii Ashukha, Dmitry Vetrov
- #80 Toward Controlled Generation of Text**
Zhiting Hu, Zichao Yang, Xiaodan Liang, Russ Salakhutdinov, Eric Xing
- #81 Robust Submodular Maximization: A Non-Uniform Partitioning Approach**
Ilija Bogunovic, Boba Mitrovic, Jonathan Scarlett, Volkan Cevher
- #82 Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning**
Jakob Foerster, Nantas Nardelli, Greg Farquhar, Triantafyllos Afouras, Phil Torr, Pushmeet Kohli, Shimon Whiteson
- #83 Stochastic Gradient Monomial Gamma Sampler**
Yizhe Zhang, Changyou Chen, Zhe Gan, Ricardo Henao, Lawrence Carin
- #84 Cost-Optimal Learning of Causal Graphs**
Murat Kocaoglu, Alex Dimakis, Sriram Vishwanath
- #85 Algebraic Variety Models for High-Rank Matrix Completion**
Greg Ongie, Laura Balzano, Rebecca Willett, Robert Nowak
- #86 Differentially Private Clustering in High-Dimensional Euclidean Spaces**
Nina Balcan, Travis Dick, Yingyu Liang, Wenlong Mou, Hongyang Zhang
- #87 Coherent probabilistic forecasts for hierarchical time series**
Souhaib Ben Taieb, James Taylor, Rob Hyndman
- #88 Unimodal Probability Distributions for Deep Ordinal Classification**
Christopher Beckham, Christopher Pal
- #89 Adversarial Feature Matching for Text Generation**
Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, Lawrence Carin
- #90 Probabilistic Submodular Maximization in Sub-Linear Time**
Serban A Stan, Morteza Zadimoghaddam, Andreas Krause, Amin Karbasi
- #91 The Predictron: End-To-End Learning and Planning**
David Silver, Hado van Hasselt, Matteo Hessel, Tom Schaul, Arthur Guez, Tim Harley, Gabriel Dulac-Arnold, David Reichert, Neil Rabinowitz, Andre Barreto, Thomas Degris
- #92 Stochastic Gradient MCMC Methods for Hidden Markov Models**
Yi-An Ma, Nick J Foti, Emily Fox
- #93 Identification and Model Testing in Linear Structural Equation Models using Auxiliary Variables**
Bryant Chen, Daniel Kumor, Elias Bareinboim
- #94 High-Dimensional Variance-Reduced Stochastic Gradient Expectation-Maximization Algorithm**
Rongda Zhu, Lingxiao Wang, Chengxiang Zhai, Quanquan Gu
- #95 Differentially Private Chi-squared Test by Unit Circle Mechanism**
Kazuya Kakizaki, Kazuto Fukuchi, Jun Sakuma
- #96 Soft-DTW: a Differentiable Loss Function for Time-Series**
Marco Cuturi, Mathieu Blondel
- #97 Learning Continuous Semantic Representations of Symbolic Expressions**
Miltos Allamanis, pankajan Chanthirasegaran, Pushmeet Kohli, Charles Sutton
- #98 On Approximation Guarantees for Greedy Low Rank Optimization**
RAJIV KHANNA, Ethan Elenberg, Alex Dimakis, Joydeep Ghosh, Sahand Negahban
- #99 Averaged-DQN: Variance Reduction and Stabilization for Deep Reinforcement Learning**
Oron Anshel, Nir Baram, Nahum Shimkin
- #100 Deep Latent Dirichlet Allocation with Topic-Layer-Adaptive Stochastic Gradient Riemannian MCMC**
Yulai Cong, Bo Chen, Hongwei Liu, Mingyuan Zhou
- #101 Estimating individual treatment effect: generalization bounds and algorithms**
Uri Shalit, Fredrik D Johansson, David Sontag
- #102 Collect at Once, Use Effectively: Making Non-interactive Locally Private Learning Possible**
Kai Zheng, Wenlong Mou, Liwei Wang
- #103 Variational Policy for Guiding Point Processes**
Yichen Wang, Grady Williams, Evangelos Theodorou, Le Song
- #104 Dance Dance Convolution**
Chris Donahue, Zachary Lipton, Julian McAuley
- #105 Language Modeling with Gated Convolutional Networks**
Yann Dauphin, Angela Fan, Michael Auli, David Grangier
- #106 Deletion-Robust Submodular Maximization: Data Summarization with “the Right to be Forgotten”**
Baharan Mirzasoleiman, Amin Karbasi, Andreas Krause
- #107 FeUdal Networks for Hierarchical Reinforcement Learning**
Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, Koray Kavukcuoglu
- #108 Distributed Batch Gaussian Process Optimization**
Erik Daxberger, Bryan Kian Hsiang Low
- #109 Recursive Partitioning for Personalization using Observational Data**
Nathan Kallus



- #110 **Optimal Densification for Fast and Accurate Minwise Hashing**
Anshumali Shrivastava
- #111 **An Adaptive Test of Independence with Analytic Kernel Embeddings**
Wittawat Jitkrittum, Zoltan Szabo, Arthur Gretton
- #112 **Deep Value Networks Learn to Evaluate and Iteratively Refine Structured Outputs**
Gygli Gygli, Mohammad Norouzi, Anelia Angelova
- #113 **World of Bits: An Open-Domain Platform for Web-Based Agents**
Tim Shi, Andrej Karpathy, Jim Fan Fan, Jonathan Hernandez, Percy Liang
- #114 **Convolutional Sequence to Sequence Learning**
Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, Yann Dauphin
- #115 **Analysis and Optimization of Graph Decompositions by Lifted Multicuts**
Andrea Hornakova, Jan-Hendrik Lange, Bjoern Andres
- #116 **Deciding How to Decide: Dynamic Routing in Artificial Neural Networks**
Mason McGill, Pietro Perona
- #117 **Scalable Multi-Class Gaussian Process Classification using Expectation Propagation**
Carlos Villacampa-Calvo, Daniel Hernandez-Lobato
- #118 **Identifying Best Interventions through Online Importance Sampling**
Rajat Sen, Karthikeyan Shanmugam, Alex Dimakis, Sanjay Shakkottai
- #119 **Stochastic Generative Hashing**
Bo Dai, Ruiqi Guo, Sanjiv Kumar, Niao He, Le Song
- #120 **Sliced Wasserstein Kernel for Persistence Diagrams**
Mathieu Carrière, Marco Cuturi, Steve Oudot
- #121 **Deeply AggreVaTeD: Differentiable Imitation Learning for Sequential Prediction**
Wen Sun, Arun Venkatraman, Geoff Gordon, Byron Boots, Drew Bagnell
- #122 **Real-Time Adaptive Image Compression**
Oren Rippel, Lubomir Bourdev
- #123 **Improved Variational Autoencoders for Text Modeling using Dilated Convolutions**
Zichao Yang, Zhiting Hu, Russ Salakhutdinov, Taylor Berg-Kirkpatrick
- #124 **Near-Optimal Design of Experiments via Regret Minimization**
Zeyuan Allen-Zhu, Yuanzhi Li, Aarti Singh, Yining Wang
- #125 **Neural Episodic Control**
Alexander Pritzel, Benigno Uria, Raam Sriram, Adrià Puigdomenech Badia, Oriol Vinyals, Demis Hassabis, Daan Wierstra, Charles Blundell
- #126 **Random Feature Expansions for Deep Gaussian Processes**
Kurt Cutajar, Edwin Bonilla, Pietro Michiardi, Maurizio Filippone
- #127 **Deep IV: A Flexible Approach for Counterfactual Prediction**
Jason Hartford, Greg Lewis, Kevin Leyton-Brown, Matt Taddy
- #128 **ZipML: Training Linear Models with End-to-End Low Precision, and a Little Bit of Deep Learning**
Hantian Zhang, Jerry Li, Kaan Kara, Dan Alistarh, Ji Liu, Ce Zhang
- #129 **Adapting Kernel Representations Online Using Submodular Maximization**
Matthew Schlegel, Yangchen Pan, Jiecao (Jack) Chen, Martha White
- #130 **End-to-End Learning for Structured Prediction Energy Networks**
David Belanger, Bishan Yang, Andrew McCallum
- #131 **Neural Message Passing for Quantum Chemistry**
Justin Gilmer, Samuel Schoenholz, Patrick F Riley, Oriol Vinyals, George Dahl
- #132 **Grammar Variational Autoencoder**
Matt J. Kusner, Brooks Paige, Jose Hernandez-Lobato
- #133 **Robust Budget Allocation via Continuous Submodular Functions**
Matthew J Staib, Stefanie Jegelka
- #134 **Neural Optimizer Search using Reinforcement Learning**
Irwan Bello, Barret Zoph, Vijay Vasudevan, Quoc Le
- #135 **Asynchronous Distributed Variational Gaussian Processes for Regression**
Hao Peng, Shandian Zhe, Cosmo Zhang, Alan Qi
- #136 **Counterfactual Data-Fusion for Online Reinforcement Learners**
Andrew Forney, Judea Pearl, Elias Bareinboim
- #137 **Large-Scale Evolution of Image Classifiers**
Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc Le, Alex Kurakin
- #138 **Spherical Structured Feature Maps for Kernel Approximation**
Yueming LYU
- #139 **A Unified View of Multi-Label Performance Measures**
Xi-Zhu Wu, Zhi-Hua Zhou
- #140 **Accelerating Eulerian Fluid Simulation With Convolutional Networks**
Jonathan Tompson, Kristofer D Schlachter, Pablo Sprechmann, Ken Perlin
- #141 **Rule-Enhanced Penalized Regression by Column Generation using Rectangular Maximum Agreement**
Jonathan Eckstein, Noam Goldberg, Ai Kagawa
- #142 **High Dimensional Bayesian Optimization with Elastic Gaussian Process**
Santu Rana, Cheng Li, Sunil Gupta, Vu Nguyen, Svetha Venkatesh
- #143 **Nyström Method with Kernel K-means++ Samples as Landmarks**
Dino Oglis, Thomas Gaertner
- #144 **Scalable Generative Models for Multi-label Learning with Missing Labels**
Vikas Jain, Nirbhay Modhe, Piyush Rai

Workshops

AUGUST 10TH & 11TH



Thursday

- **Workshop on Computational Biology**
Level 4, C4.4
- **Lifelong Learning: A Reinforcement Learning Approach**
Level 4, C4.1
- **Automatic Machine Learning (AutoML 2017)**
Level 4, C4.9
- **ICML Workshop on Machine Learning for Autonomous Vehicles 2017**
Level 4, C4.10
- **Implicit Generative Models**
Level 2, Parkside 1
- **Learning to Generate Natural Language**
Level 4, C4.11
- **Principled Approaches to Deep Learning**
Level 4, C4.5
- **Video Games and Machine Learning**
Level 4, C4.6
- **Workshop on Human Interpretability in Machine Learning (WHI)**
Level 4, C4.8
- **Workshop on Visualization for Deep Learning**
Level 4, C4.3
- **ML on a budget: IoT, Mobile and other tiny-ML applications**
Level 4, C4.7

Friday

- **Human in the Loop Machine Learning**
Level 4, C4.8
- **Deep Structured Prediction**
Level 4, C4.5
- **Reproducibility in Machine Learning Research**
Level 4, C4.10
- **Machine Learning in Speech and Language Processing**
Level 4, C4.3
- **Reliable Machine Learning in the Wild**
Level 4, C4.7
- **Picky Learners: Choosing Alternative Ways to Process Data**
Level 4, C4.6
- **Machine Learning for Music Discovery**
Level 4, C4.9
- **Time Series Workshop**
Level 4, C4.1
- **Reinforcement Learning Workshop**
Level 2, Parkside 1
- **Interactive Machine Learning and Semantic Information Retrieval**
Level 4, C4.11
- **Private and Secure Machine Learning**
Level 4, C4.4

Thursday

AUGUST 10TH | WORKSHOPS



THURSDAY WORKSHOPS 8:30 AM - 5:30 PM

Workshop on Computational Biology

Location: Level 4, C4.4

<i>Dana Pe'er</i>	<i>MSKCC</i>
<i>Christina Leslie</i>	<i>MSKCC</i>
<i>Barbara Engelhardt</i>	<i>Princeton University</i>
<i>Elham Azizi</i>	<i>MSKCC</i>
<i>Sandhya Prabhakaran</i>	<i>MSKCC</i>
<i>Meghana Kshirsagar</i>	<i>MSKCC</i>
<i>Ambrose Carr</i>	<i>MSKCC</i>

The workshop will showcase recent research in the field of Computational Biology. There has been significant development in genomic sequencing techniques as well as imaging technologies that not only generate huge amounts of data but provide unprecedented levels of resolution, that of a single cell and even subcellular resolution. This availability of high dimensional data, at multiple spatial and temporal resolutions and capturing several perspectives of biological phenomena has made machine learning methods increasingly relevant for computational analysis of the data. Conversely, biological data has also exposed unique challenges and problems that call for the development of new machine learning methods. This workshop aims at bringing in researchers working at the intersection of Machine Learning and Biology to present recent advances and open questions in computational biology to the ICML community.

<https://sites.google.com/view/compbioworkshopicml2017>

Lifelong Learning: A Reinforcement Learning Approach

Location: Level 4, C4.1

<i>Sarath Chandar</i>	<i>University of Montreal</i>
<i>Balaraman Ravindran</i>	<i>Indian Inst. of Tech Madras</i>
<i>Daniel J. Mankowitz</i>	<i>Technion Israel Inst. of Tech</i>
<i>Shie Mannor</i>	<i>Technion Israel Inst. of Tech</i>
<i>Tom Zahavy</i>	<i>Technion Israel Inst. of Tech</i>

One of the most challenging and open problems in Artificial Intelligence (AI) is that of Lifelong Learning:

“Lifelong Learning is the continued learning of tasks, from one or more domains, over the course of a lifetime, by a lifelong learning system. A lifelong learning system efficiently and effectively (1) retains the knowledge it has learned; (2) selectively transfers knowledge to learn new tasks; and (3) ensures the effective and efficient interaction between (1) and (2).”.

Lifelong learning is still in its infancy. Many issues currently exist such as learning general representations, catastrophic forgetting, efficient knowledge retention mechanisms and hierarchical abstractions.

<http://rlabstraction2016.wixsite.com/icml-2017>



Automatic Machine Learning (AutoML 2017)

Location: Level 4, C4.9

Joaquin Vanschoren *TU Eindhoven / OpenML*
Roman Garnett *Washington University in St. Louis*

Machine learning has achieved considerable successes in recent years and an ever-growing number of disciplines rely on it. However, this success crucially relies on human machine learning experts, who select appropriate features, workflows, machine learning paradigms, algorithms, and their hyperparameters. As the complexity of these tasks is often beyond non-experts, the rapid growth of machine learning applications has created a demand for off-the-shelf machine learning methods that can be used easily and without expert knowledge. We call the resulting research area that targets progressive automation of machine learning AutoML.

<https://sites.google.com/site/automl2017icml/>

ICML Workshop on Machine Learning for Autonomous Vehicles 2017

Location: Level 4, C4.10

Li Erran Li *Uber Technologies*
Raquel Urtasun *University of Toronto*
Andrew Gray *Uber Technologies*
Silvio Savarese *Stanford University*

Although dramatic progress has been made in the field of autonomous driving, there are many major challenges in achieving full-autonomy. For example, how to make perception accurate and robust to accomplish safe autonomous driving? How to reliably track cars, pedestrians, and cyclists? How to learn long term driving strategies (known as driving policies) so that autonomous vehicles can be equipped with adaptive human negotiation skills when merging, overtaking and giving way, etc? How to achieve near-zero fatality?

<https://sites.google.com/site/ml4autovehicles2017>

Implicit Generative Models

Location: Level 2, Parkside 1

Rajesh Ranganath *Princeton University*
Ian Goodfellow *Google Brain*
Dustin Tran *Columbia University*
David Blei *Columbia University*
Balaji Lakshminarayanan *Google DeepMind*
Shakir Mohamed *DeepMind*

Probabilistic models are a central implement in machine learning practice. They form the basis for models that generate realistic data, uncover hidden structure, and make predictions. Traditionally, probabilistic models in machine learning have focused on prescribed models. Prescribed models specify a joint density over observed and hidden variables that can be easily evaluated. The requirement of a tractable density simplifies their learning but limits their flexibility --- several real world phenomena are better described by simulators that do not admit a tractable density. Probabilistic models defined only via the simulations they produce are called implicit models.

<https://sites.google.com/view/implicitmodels>

Learning to Generate Natural Language

Location: Level 4, C4.11

Phil Blunsom *(DeepMind and University of Oxford)*
Kris Cao *(University of Cambridge)*
Stephen Clark *(DeepMind and University of Cambridge)*
Chris Dyer *(DeepMind and Carnegie Mellon University)*
Milica Gašić *(University of Cambridge)*
Daniela Gerz *(University of Cambridge)*
Wang Ling *(DeepMind)*
Yishu Miao *(University of Oxford)*
Tsung-Hsien Wen *(University of Cambridge)*

Research on natural language generation is rapidly growing due to the increasing demand for human-machine communication in natural language. This workshop aims to promote the discussion, exchange, and dissemination of ideas on the topic of text generation, touching several important aspects in this modality: learning schemes and evaluation, model design and structures, advanced decoding strategies, and natural language generation applications. This workshop aims to be a venue for the exchange of ideas regarding data-driven machine learning approaches for text generation, including mainstream tasks such



as dialogue generation, instruction generation, and summarization; and for establishing new directions and ideas with potential for impact in the fields of machine learning, deep learning, and NLP.

<https://sites.google.com/site/langgen17/home>

Principled Approaches to Deep Learning

Location: Level 4, C4.5

<i>Andrzej Pronobis</i>	<i>University of Washington</i>
<i>Robert Gens</i>	<i>Google</i>
<i>Sham M. Kakade</i>	<i>University of Washington</i>
<i>Pedro Domingos</i>	<i>University of Washington</i>

The recent advancements in deep learning have revolutionized the field of machine learning, enabling unparalleled performance and many new real-world applications. Yet, the developments that led to this success have often been driven by empirical studies, and little is known about the theory behind some of the most successful approaches. While theoretically well-founded deep learning architectures had been proposed in the past, they came at a price of increased complexity and reduced tractability. Recently, we have witnessed considerable interest in principled deep learning. This led to a better theoretical understanding of existing architectures as well as development of more mature deep models with solid theoretical foundations. In this workshop, we intend to review the state of those developments and provide a platform for the exchange of ideas between the theoreticians and the practitioners of the growing deep learning community. Through a series of invited talks by the experts in the field, contributed presentations, and an interactive panel discussion, the workshop will cover recent theoretical developments, provide an overview of promising and mature architectures, highlight their challenges and unique benefits, and present the most exciting recent results.

<http://padl.ws/>

Video Games and Machine Learning

Location: Level 4, C4.6

<i>Gabriel Synnaeve</i>	<i>Facebook AI Research</i>
<i>Julian Togelius</i>	<i>NYU</i>
<i>Tom Schaul</i>	<i>Google DeepMind</i>
<i>Oriol Vinyals</i>	<i>Google DeepMind</i>
<i>Nicolas Usunier</i>	<i>Facebook AI Research</i>

This workshop focuses on complex games which provide interesting and hard challenges for machine learning. Going beyond simple toy problems of the past, and games which can easily be solved with search, we focus on games where learning is likely to be necessary to play well. This includes strategy games such as StarCraft [4,5], open-world games such as MineCraft [6,7,8], first-person shooters such as Doom [9,10], as well as hard and unsolved 2D games such as Ms. Pac-Man and Montezuma’s Revenge [11,12,13]. While we see most of the challenges in game-playing, there are also interesting machine learning challenges in modeling and content generation [14]. This workshop aims at bringing together all researchers from ICML who want to use video games as a benchmark. We will have talks by invited speakers from machine learning, from the game AI community, and from the video games industry.

https://syhw.github.io/vgml_workshop_icml2017

Workshop on Human Interpretability in Machine Learning (WHI)

Location: Level 4, C4.8

<i>Been Kim</i>	<i>Google Brain</i>
<i>Dmitry Malioutov</i>	<i>The D. E. Shaw Group</i>
<i>Kush R. Varshney</i>	<i>IBM Research</i>
<i>Adrian Weller</i>	<i>University of Cambridge</i>

This workshop will bring together researchers who study the interpretability of predictive models, develop interpretable machine learning algorithms, and develop methodology to interpret black-box machine learning models (e.g., post-hoc interpretations). This is a very exciting time to study interpretable machine learning, as the advances in large-scale optimization and Bayesian inference that have enabled the rise of



black-box machine learning are now also starting to be exploited to develop principled approaches to large-scale interpretable machine learning. Participants in the workshop will exchange ideas on these and allied topics.

<https://sites.google.com/view/whi2017/home>

Workshop on Visualization for Deep Learning

Location: Level 4, C4.3

<i>John Canny</i>	<i>UC Berkeley</i>
<i>Polo Chau</i>	<i>Georgia Tech</i>
<i>Xiangmin Fan</i>	<i>Chinese Academy of Sciences</i>
<i>Biye Jiang</i>	<i>UC Berkeley</i>
<i>Jun-Yan Zhu</i>	<i>UC Berkeley</i>

Deep networks have had profound impact across machine learning research and in many application areas. DNNs are complex to design and train. They are non-linear systems that almost always have many local optima and are often sensitive to training parameter settings and initial state. Systematic optimization of structure and hyperparameters is possible e.g. with Bayesian optimization, but hampered by the expense of training each design on realistic datasets. Exploration is still ongoing for best design principles. We argue that visualization can play an essential role in understanding DNNs and in developing new design principles. With rich tools for visual exploration of networks during training and inference, one should be able to form closer ties between theory and practice: validating expected behaviors, and exposing the unexpected which can lead to new insights. With the rise of generative modeling and reinforcement learning, more interesting directions like understanding and visualization of generative models, visual explanation for driving policy could be explored as well.

As the second edition of this workshop, we are proposing changes based on the lessons we learned last year. We would like to organize a few domain specific tutorials, and panel discussions. We do think machine learning researchers need a lot of tutorials

and advice from the visualization/HCI community and vice versa. Many audience in our workshop last year also suggested that more discussion can greatly help us better define such interdisciplinary area.

<http://icmlviz.github.io/>

ML on a budget: IoT, Mobile and other tiny-ML applications

Location: Level 4, C4.7

<i>Manik Varma</i>	<i>Microsoft Research</i>
<i>Vankatesh Saligrama</i>	<i>Boston University</i>
<i>Prateek Jain</i>	<i>Microsoft Research</i>

We routinely encounter scenarios where at test-time we must predict on a budget. Feature costs in Internet, Healthcare, and Surveillance applications arise due to feature extraction time and feature/sensor acquisition~\cite{trapeznikov:2013b} costs. Data analytics applications in mobile devices are often performed on remote cloud services due to the limited device capabilities, which imposes memory/prediction time costs. Naturally, in these settings, one needs to carefully understand the trade-off between accuracy and prediction cost. Uncertainty in the observations, which is typical in such scenarios, further adds to complexity of the task and requires a careful understanding of both the uncertainty as well as accuracy-cost tradeoffs.

In this workshop, we aim to bring together researchers from various domains to discuss the key aspects of the above mentioned emerging and critical topic. The goal is to provide a platform where ML/statistics/optimization researchers can interact closely with domain experts who need to deploy ML models in resource-constrained settings (like an IoT device maker), and chart out the foundational problems in the area and key tools that can be used to solve them.

<https://sites.google.com/site/tinyml2017/>

Friday

AUGUST 11TH | WORKSHOPS



FRIDAY WORKSHOPS 8:30 AM - 5:30 PM

Human in the Loop Machine Learning

Location: Level 4, C4.8

Richard Nock

*Data61, Australian National Univ.
University of Sydney*

Cheng Soon Ong

Data61, Australian National Univ.

As machine learning systems become more ubiquitous in everybody's day-to-day life or work, society and industry is in an intermediate state between fully manual and fully automatic systems. The gradient undoubtedly points towards full automation, but moving forward in this direction is going to face increasing challenges due to the fact that current machine learning research tends to focus on end to end systems, which puts aside the fact that for practical applications there are still gaps or caveats in the automation. Parts of these come from the presence of (or the necessity to have) the Human in the Loop. There are two main locations for the Human in the automated system: (i) upstream, in which case the focus is mainly in the inputs of the algorithm. (ii) downstream: other domains have machine learning approaches analyse parts of the data, and human experts use the results and intuition to make decisions.

This workshop aims to bring together people who are working on systems where machine learning is only part of the solution. Participants will exchange ideas and experiences on human in the loop machine learning.

<http://machlearn.gitlab.io/hitl2017/>

Deep Structured Prediction

Location: Level 4, C4.5

Isabelle Augenstein

University of Copenhagen

Kai-Wei Chang

UCLA

Gal Chechik

Bar-Ilan University / Google

Bert Huang

Virginia Tech

Andre Martins

Unbabel & Inst. de Telecomun.

Ofer Meshi

Google

Yishu Miao

University of Oxford

Alexander Schwing

University of Illinois Urbana

In recent years, deep learning has revolutionized machine learning. Most successful applications of deep learning involve predicting single variables (e.g., univariate regression or multi-class classification). However, many real problems involve highly dependent, structured variables. In such scenarios, it is desired or even necessary to model correlations and dependencies between the multiple input and output variables. Such problems arise in a wide range of domains, from natural language processing, computer vision, computational biology and others. This workshop will bring together experts in machine learning and application domains whose research focuses on combining deep learning and structured models. Specifically, we aim to provide an overview of existing approaches from various domains to distill from their success principles that can be more generally applicable. We will also discuss the main challenges that arise in this setting and outline potential directions for future progress. The target audience consists of researchers and practitioners in machine learning and application areas.

<https://deepstruct.github.io/ICML17>



Reproducibility in Machine Learning Research

Location: Level 4, C4.10

<i>Anirudh Goyal</i>	<i>(MILA) Université de Montréal</i>
<i>Rosemary Nan Ke</i>	<i>(MILA) Université de Montréal</i>
<i>Alex Lamb</i>	<i>(MILA) Université de Montréal</i>
<i>Joelle Pineau</i>	<i>Mcgill University</i>
<i>Samy Bengio</i>	<i>Google Brain</i>
<i>Yoshua Bengio</i>	<i>Université de Montréal</i>

This workshop focuses on issues of reproducibility and replication of results in the Machine Learning community. Papers from the Machine Learning community are supposed to be a valuable asset. They can help to inform and inspire future research. They can be a useful educational tool for students. They can give guidance to applied researchers in industry. Perhaps most importantly, they can help us to answer the most fundamental questions about our existence - what does it mean to learn and what does it mean to be human? Reproducibility, while not always possible in science (consider the study of a transient astrological phenomenon like a passing comet), is a powerful criteria for improving the quality of research. A result which is reproducible is more likely to be robust and meaningful and rules out many types of experimenter error (either fraud or accidental).

<https://sites.google.com/view/icml-reproducibility-workshop>

Machine Learning in Speech and Language Processing

Location: Level 4, C4.3

<i>Karen Livescu</i>	<i>TTI-Chicago</i>
<i>Anton Ragni</i>	<i>U. Cambridge</i>
<i>Liang Lu</i>	<i>TTI-Chicago</i>
<i>Tara Sainath</i>	<i>Google</i>

This workshop continues a tradition of MLSLP workshops held as satellites of ICML, ACL, and Interspeech conferences. While research in speech and language processing has always involved machine learning (ML), current research is benefiting from even closer interaction between these fields. Speech and language processing is continually mining new ideas from ML and ML, in turn, is devoting more interest to speech and language applications. This workshop is a venue for locating and incubating the

next waves of research directions for interaction and collaboration. The workshop will (1) discuss emerging research ideas with potential for impact in speech/language and (2) bring together relevant researchers from ML and speech/language who may not regularly interact at conferences. Example topics include new directions for deep learning in speech/language, reinforcement learning, unsupervised/semi-supervised learning, domain adaptation/transfer learning, and topics at the boundary of speech, text, and other modalities.

<http://ttic.uchicago.edu/~klivescu/MLSLLP2017>

Reliable Machine Learning in the Wild

Location: Level 4, C4.7

<i>Dylan Hadfield-Menell</i>	<i>UC Berkeley</i>
<i>Jacob Steinhardt</i>	<i>Stanford University</i>
<i>Adrian Weller</i>	<i>University of Cambridge</i>
<i>Smitha Milli</i>	<i>OpenAI, UC Berkeley</i>

When can we trust that a system that has performed well in the past will continue to do so in the future? Designing systems that are reliable in the wild is essential for high stakes applications such as self-driving cars and automated surgical assistants. This workshop aims to bring together researchers in diverse areas such as reinforcement learning, human-robot interaction, game theory, cognitive science, and security to further the field of reliability in machine learning. We will focus on three aspects — robustness (to adversaries, distributional shift, model misspecification, corrupted data); awareness (of when a change has occurred, when the model might be miscalibrated, etc.); and adaptation (to new situations or objectives).

We aim to consider each of these in the context of the complex human factors that impact the successful application or meaningful monitoring of any artificial intelligence technology. Together, these will aid us in designing and deploying reliable machine learning systems.

<https://sites.google.com/site/wildml2017icml>



Picky Learners: Choosing Alternative Ways to Process Data

Location: Level 4, C4.6

<i>Kamalika Chaudhuri</i>	<i>UC San Diego</i>
<i>Corinna Cortes</i>	<i>Google Research</i>
<i>Giulia DeSalvo</i>	<i>New York University</i>
<i>Chicheng Zhang</i>	<i>UC San Diego</i>
<i>Ningshan Zhang</i>	<i>New York University</i>

Picky Learners consists of a broad range of learning scenarios where the learner does not simply process every data point blindly, but instead can choose to incorporate them in alternative ways. Despite the growing costs of processing and labelling vast amounts of data, only isolated efforts have tackled this problem primarily in the areas of active learning, learning with rejection and on-line learning with feedback graphs.

The goal of this workshop is then to bring together researchers and practitioners in these three areas in order to bridge the gap between active learning, learning with rejection, and on-line learning with feedback graphs. We expect that the fruitful collaborations started in this workshop will result in novel research that will help develop each field.

<https://sites.google.com/site/rejectionactiveicml>

Machine Learning for Music Discovery

Location: Level 4, C4.9

<i>Erik Schmidt</i>	<i>Pandora</i>
---------------------	----------------

The ever-increasing size and accessibility of vast music libraries has created a demand more than ever for machine learning systems that are capable of understanding and organizing this complex data. While this topic has received relatively little attention within the machine learning community, it has been an area of intense focus within the community of Music Information Retrieval (MIR), where significant progress has been made, but these problems remain far from solved. Furthermore, the recommender systems community has made great progress in terms of collaborative feedback recommenders, but these approaches suffer strongly from the cold-start problem. As such, recommendation techniques often fall back on content-based machine learning systems, but defining

musical similarity is extremely challenging as myriad features all play some role (e.g., cultural, emotional, timbral, rhythmic).

<https://sites.google.com/site/ml4md2017/>

Time Series Workshop

Location: Level 4, C4.1

<i>Vitaly Kuznetsov</i>	<i>Google Research</i>
<i>Yan Liu</i>	<i>University of Southern California</i>
<i>Scott Yang</i>	<i>Courant Institute</i>
<i>Rose Yu</i>	<i>University of Southern California</i>

Time series data is ubiquitous. In domains as diverse as finance, entertainment, transportation and health-care, there has been a fundamental shift away from parsimonious, infrequent measurement to nearly continuous monitoring and recording. Rapid advances in many sensing technologies, ranging from remote sensors to wearables and social sensing, are have generated a rapid growth in the size and complexity of time series archives. Thus, while time series analysis has been studied extensively in the past, its importance only continues to grow. Furthermore, modern time series data pose significant new challenges in terms of structure (e.g., irregular sampling in hospital records and spatiotemporal structure in climate data) and size (e.g. computation and storage). These challenges are compounded by the fact that the standard i.i.d. assumptions used in other areas of machine learning are often not appropriate for time series. Instead, new theory, models and algorithms are needed to process and analyze this data.

<http://roseyu.com/time-series-workshop/>



Reinforcement Learning Workshop

Location: Level 2, Parkside 1

Doina Precup

McGill University

The workshop will contain presentations of late-breaking reinforcement learning results in all areas of the field, including deep reinforcement learning, exploration, transfer learning and using auxiliary tasks, theoretical result etc, as well as applications of reinforcement learning to various domains. A panel discussion on the most interesting and challenging current research directions will conclude the workshop.

Interactive Machine Learning and Semantic Information Retrieval

Location: Level 4, C4.11

Dorota Glowacka

University of Helsinki

Wray Buntine

Monash University

Petri Myllymaki

University of Helsinki

Retrieval techniques operating on text or semantic annotations have become the industry standard for retrieval from large document collections. However, traditional information retrieval techniques operate on the assumption that the user issues a single query and the system responds with a ranked list of documents. In recent years we have witnessed a substantial growth in text data coming from various online resources, such as online newspapers, blogs, specialized document collections (e.g. arXiv). Traditional information retrieval approaches often fail to provide users with adequate support when browsing such online resources, hence in recent years there has been a growing interest in developing new algorithms and design methods that can support interactive information retrieval. The aim of this workshop is to explore new methods and related system design for interactive data analytics and management in various domains, including specialized text collections (e.g. legal, medical, scientific) as well as for various tasks, such as semantic information retrieval, conceptual organization and clustering of data collections for sense making, semantic expert profiling, and document recommender systems. Of interest, also, is probabilistic and machine learning formulations of the interactive information retrieval task above and beyond the simple “language models” framework developed in the information retrieval community.

<https://sites.google.com/view/interactive-ml>

Private and Secure Machine Learning

Location: Level 4, C4.4

Antti Honkela

University of Helsinki

Kana Shimizu

Waseda University

Samuel Kaski

Aalto University

There are two complementary approaches to private and secure machine learning: differential privacy can guarantee privacy of the subjects of the training data with respect to the output of a differentially private learning algorithm, while cryptographic approaches can guarantee secure operation of the learning process in a potentially distributed environment. The aim of this workshop is to bring together researchers interested in private and secure machine learning, to stimulate interactions to advance either perspective or to combine them.

<https://sites.google.com/view/psml>



Funding for our Sponsor Scholars program was generously provided by our platinum sponsors, American Express and Baidu, as well as the National Science Foundation and the Artificial Intelligence Journal. Their exemplary support provided scholarships to help 217 student researchers travel to ICML to present their research. We gratefully acknowledge these sponsors, and congratulate all of our Sponsor Scholars, listed below.

- Massil Achab**, Ecole Polytechnique
Olanike Akinduyite, Federal U. of Tech., Akure
Sadura Akinrinwa, Federal U. of Tech., Akure
Riad Akrouf, TU Darmstadt
Brandon Amos, CMU
Ash Anand, Indian Inst. of Tech., Kharagpur
Kavosh Asadi, Brown U.
Arsenii Ashukha, HSE, MIPT
Matej Balog, Cambridge and MPI Tübingen
Ekaterina N Beresneva, NRU HSE
Garrett Bernstein, UMass
Jonathan Binas, MILA, U. of Montreal
Umanga Bista, Australian National U.
Sean Bittner, Columbia U.
Ashish Bora, U. of Texas
Alex Botev, UCL
Ahcène Boubekki, Leuphana
Matko Bošnjak, UCL
Noam Brown, CMU
Victor Brunel, MIT
Thang Bui, Cambridge
Daniele Calandriello, INRIA Lille
Mathieu Carrière, Inria Saclay
Sarath Chandar, U. of Montreal
Weiling Chen, Nanyang Technological U.
Jiecao (Jack) Chen, Indiana U. Bloomington
Sheng Chen, U. of Minnesota
Yichen Chen, Princeton U.
Xixian Chen, The Chinese U. of Hong Kong
Bangrui Chen, Cornell U.
Christopher Chow, Australian National U.
Yulai Cong, Xidian U.
Lidia Contreras-Ochando, U. Politècnica de Valencia
Kurt Cutajar, EURECOM
Carlo D'Eramo, Politecnico di Milano
Bo Dai, Georgia Tech
Erik Daxberger, Ludwig-Maximilians-U. München
Giulia DeSalvo, New York U.
Walter Dempsey, U. of Michigan
Yuntian Deng, Harvard U.
Natalia Diaz Rodriguez, ENSTA ParisTech
Chris Donahue, UC San Diego
- Simon Du**, CMU
Ethan Elenberg, The U. of Texas at Austin
Adel Elmahdy, U. of Minnesota
Rahim Entezari, Iran U. of Science and Tech.
Rasool Fakoore, U. of Texas at Arlington
Gabriele Farina, CMU
Jean Feng, U. of Washington
Cristina I. Font, Polytechnic U. of Valencia
Andrew Forney, UCLA
Luca Franceschi, Inst. Italiano di Tec. - UCL
Joe Futoma, Duke U.
Wenbo Gao, Columbia U.
Guillaume Gautier, INRIA Lille
Dar Gilboa, Columbia
Insu Han, Korea Advanced Inst. of Science and Tech.
Josiah Hanna, U. of Texas
Sanjay Hariharan, Duke U.
Jason Hartford, U. of British Columbia
Bryan He, Stanford
Lifang He, U. of Illinois at Chicago/Shenzhen U.
Kyle Helfrich, U. of Kentucky
Oliver Hinder, Stanford
Shuguang Hu, U. of Hong Kong
Weihua Hu, The U. of Tokyo / RIKEN
Zhiting Hu, CMU
Ajil Jalal, U. of Texas
Neal Jean, Stanford U.
Yacine Jernite, New York U.
Shali Jiang, Washington U. in St. Louis
Li Jing, MIT
Wittawat Jitkittum, Gatsby Unit, UCL
Matthew Joseph, U. of Pennsylvania
Gautam Kamath, MIT
Kirthevasan Kandasamy, CMU
Kaan Kara, ETH Zurich
Sammie Katt, Northeastern
Ambika Kaul, Intl. Inst. of Information Tech., Hyderabad
Murat Kocaoğlu, U. of Texas
Pang Wei Koh, Stanford
Alexander Kolesnikov, IST Austria
Daniel Kumor, Purdue U.
Hima Lakkaraju, Stanford
Juho Lee, POSTECH
Kimin Lee, KAIST
- Qi Lei**, U. of Texas at Austin
Hao Li, U. of Maryland at College Park
Wei-Chen Li, National Taiwan U.
Shuai Li, U. of Cambridge
Hanbo Li, UC San Diego
Yingzhen Li, Cambridge
Qunwei Li, Syracuse U.
Hanxiao Liu, CMU
Simon Luo, U. of Tasmania; Data61, CSIRO
Yueming Lyu, city U. of Hong Kong
Ramin M. Hasani, Vienna U. of Tech.
Yi-An Ma, U. of Washington
Saket Maheshwary, Intl. Inst. of Information Tech., Hyderabad
Sebastian Mair, Leuphana U. Lüneburg
Gustavo Malkomes, Washington U. in St. Louis
Raunak Manekar, BITS Pilani
Xueyu Mao, U. of Texas
Ium Maystre, EPFL
Rowan McAllister, Cambridge
Daniel McNamara, Australian National U. and Data61
Jiali Mei, EDF R&D & Université Paris-Sud
Yishu Miao, U. of Oxford
Andrew Miller, Harvard
Baharan Mirzasoleiman, ETH Zurich
Nikhil Mishra, UC Berkeley
Marko Mitrovic, Yale U.
Boba Mitrovic, EPFL
Søren Wengel Mogensen, U. of Copenhagen
Karim Moradi, U. of Kurdistan
Alexander Moreno, Georgia Inst. of Tech.
Wenlong Mou, Peking U.
Jonas Mueller, MIT
Tsenduren Munkhdalai, UMass
Cameron Musco, MIT
Tasha Nagamine, Columbia U.
Willie Neiswanger, CMU
Xiuyan Ni, The Graduate Center, CUNY
Lachlan J Nicholson, Queensland U. of Tech.
Dino Oglic, U. of Bonn
Junhyuk Oh, U. of Michigan
Kazeem Olorisade, Keele U.
Yunpeng Pan, Georgia Tech
- Nicolas Papernot**, Google / Penn State
Sahar Qaadani, Neural computation
Aditi Raghunathan, Stanford
Anant Raj, Max-Planck Inst. for Intelligent Systems
Sayak Ray Chowdhury, Indian Inst. of Science
Hippolyt Ritter, U. College London
Pablo Rozas Larraondo, Australian National U.
Ariel Ruiz-Garcia, Coventry U.
Tammo Rukat, U. of Oxford
Arun Suggala, CMU
Tomoya Sakai, The U. of Tokyo / RIKEN
Emanuele Sansone, U. of Trento
Daniel Selsam, Stanford
Rajat Sen, U. of Texas
Vatsal Sharan, Stanford U.
Ehsan Shareghi, Monash U.
Utkarsh Sharma, U. of Sydney
Yichen Shen, MIT
Jie Shen, Rutgers
Li Shen, Tencent
Zhan Shi, U. of Illinois at Chicago
Tim Shi, Stanford
Rui Shu, Stanford
Pranav Shyam, R. V. College of Engineering & Indian Inst. of Science
Karan Singh, Princeton U.
Shashank B, CMU
Vidyashankar Sivakumar, U. of Minnesota
Kacper Sokol, U. of Bristol
Seok-Jeong Song, Kyung Hee U.
Mari L. Soto, MIT/Harvard/CUNY
Matthew Staib, MIT
Ke Sun, KAUST
Wen Sun, CMU
Yu Sun, Cornell U.
Dougal Sutherland, Gatsby unit, U. College London
Zilong Tan, Duke U.
Rashish Tandon, U. of Texas
Xiangru Tang, Central China Normal U.
Wesley Tansey, U. of Texas at Austin
Anna Thomas, Stanford
Samuele Tosatto, Politecnico di Milano
Chris Tosh, UC San Diego
Nilesh Tripuraneni, UC Berkeley
- Rakshit Trivedi**, Georgia Inst. of Tech.
Jenkin Tsui, U. of Toronto
Shashanka Ubaru, U. of Minnesota
Jonas Umlauf, Technical U. of Munich
Sharan Vaswani, U. of British Columbia
Carlos Villacampa-Calvo, Universidad Autónoma de Madrid
Ruben Villegas, U. of Michigan
Riccardo Volpi, RIST - Romanian Inst. of Science
Lingxiao Wang, U. of Virginia
Yu-Xiang Wang, CMU / Amazon AWS
Shen Wang, U. of Illinois
Zi Wang, MIT
Yining Wang, CMU
Yunhe Wang, Peking U.
Yixin Wang, Columbia U.
Yichen Wang, Gatech
Pengfei Wei, Nanyang Technological U., Singapore
Kevin Winner, UMass
Andrew Wrigley, Australian National U.
Xi-Zhu Wu, Nanjing U.
Pengtao Xie, CMU
Pan Xu, U. of Virginia
Zheng Xu, U. of Maryland
Scott Yang, Courant Inst.
Lin Yang, Johns Hopkins U.
Hongyu Yang, MIT
Haishan Ye, Shanghai Jiao Tong U.
Jianbo Ye, Penn State U.
Manzil Zaheer, CMU
Jingfeng Zhang, National U. of Singapore
Jiong Zhang, U. of Texas
Xiao Zhang, U. of Virginia
Yilin Zhang, U. of Wisconsin-Madison
Wenpeng Zhang, Tsinghua U.
Yizhe Zhang, Duke U.
Zhezhong Zhang, Zhejiang U. & Tencent AI Lab
Chicheng Zhang, UCSD
Hantian Zhang, ETH Zurich
Teng Zhang, Nanjing U.
He Zhao, FIT, Monash U.
Shengjia Zhao, Stanford
Mingmin Zhao, MIT
Shuai Zheng, Hong Kong U. of Science and Tech.
Kai Zheng, Peking U.
Kai Zhong, U. of Texas
Hao Zhou, U. of Wisconsin - Madison



- Abbeel, Pieter:** Talk Mon in Reinforcement learning 2, Talk Mon in Deep learning 3: metalearning, Poster Mon #112, Poster Mon #61, Poster Mon #52, Poster Mon #43
- Achab, Massil:** Talk Wed in Causal Inference 1, Poster Wed #75
- Acharya, Jayadev:** Talk Tue in Information theory, Poster Tue #128
- Achiam, Joshua:** Talk Mon in Reinforcement learning 2, Poster Mon #43
- Adams, Ryan:** Talk Tue in Probabilistic inference 3, Poster Tue #109
- Afouras, Triantafyllos:** Talk Wed in Deep reinforcement learning 1, Poster Wed #82
- Agarwal, Naman:** Talk Wed in Privacy and security 2, Poster Wed #68
- Agarwal, Alekh:** Tutorial Sun in Tutorials Session A, Talk Mon in Active learning, Poster Mon #142, Talk Tue in Reinforcement learning 5, Poster Tue #49, Poster Tue #65
- Agrawal, Pulkit:** Talk Mon in Reinforcement learning 3, Poster Mon #88
- Ahmadi Fahandar, Mohsen:** Talk Tue in Ranking and preferences, Poster Tue #81
- Ahmed, Amr:** Talk Tue in Recurrent neural networks 4, Poster Tue #138, Talk Wed in Monte Carlo methods 1, Poster Wed #39
- Akrour, Riadh:** Talk Wed in Continuous control, Poster Wed #55
- Aksoylar, Cem:** Talk Mon in Continuous optimization 2, Poster Mon #78
- Aldà, Francesco:** Talk Wed in Privacy and security 2, Poster Wed #77
- Ali, Alnur:** Talk Mon in Continuous optimization 4, Poster Tue #7
- Alistarh, Dan:** Talk Wed in Large scale learning, Poster Wed #128
- Allamanis, Miltiadis:** Talk Wed in Language 2, Poster Wed #97
- Allen-Zhu, Zeyuan:** Tutorial Sun in Tutorials Session A, Talk Mon in Online learning 2, Poster Mon #63, Talk Tue in Spectral methods, Talk Tue in Continuous optimization 7, Poster Tue #115, Poster Tue #44, Poster Tue #68, Talk Wed in Combinatorial optimization 2, Poster Wed #124
- Amato, Chris:** Talk Mon in Reinforcement learning 3, Poster Mon #115, Talk Tue in Game theory and multiagents, Poster Tue #140
- Amos, Brandon:** Talk Tue in Deep learning 6, Poster Wed #5, Poster Wed #7
- Anderson, David:** Talk Mon in Matrix factorization 3, Poster Tue #25
- Andreas, Jacob:** Talk Mon in Reinforcement learning 4, Poster Tue #27
- Andres, Bjoern:** Talk Wed in Combinatorial optimization 2, Poster Wed #115
- Angelova, Anelia:** Talk Wed in Structured prediction, Poster Wed #112
- Anschel, Oron:** Talk Mon in Reinforcement learning 3, Poster Mon #106
- Anschel, Oron:** Talk Wed in Deep reinforcement learning 1, Poster Wed #99
- Appel, Ron:** Talk Tue in Ensemble methods, Poster Tue #72
- Archambeau, Cedric:** Talk Wed in Bayesian Optimization, Poster Wed #40
- Arik, Sercan:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Arjevani, Yossi:** Talk Mon in Continuous optimization 1, Poster Mon #11
- Arjovsky, Martin:** Talk Mon in Deep generative models 2, Poster Mon #77
- Arora, Sanjeev:** Talk Mon in Deep generative models 2, Poster Mon #41
- Arpit, Devansh:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Asadi, Kavosh:** Talk Mon in Reinforcement learning 2, Poster Mon #70
- Ashukha, Arsenii:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #79
- Aspuru-Guzik, alan:** Talk Wed in Bayesian Optimization, Poster Wed #57
- Atia, George:** Talk Mon in Matrix factorization 2, Poster Mon #118, Talk Tue in Sparsity 2, Poster Tue #127
- Augenstein, Isabelle:** Workshop Fri in C4.5
- Auli, Michael:** Talk Wed in Language 3, Poster Wed #114, Poster Wed #105
- Avron, Haim:** Talk Mon in Supervised learning 2, Poster Mon #66
- Azizi, Elham:** Workshop Thu in C4.4
- Bach, Stephen:** Talk Mon in Probabilistic learning 3, Poster Mon #89
- Bach, Francis:** Talk Wed in Distributed optimization, Poster Wed #37
- Bachem, Olivier:** Talk Tue in Clustering 1, Talk Tue in Learning theory 2, Poster Tue #37, Poster Tue #76
- Bachman, Philip:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #21
- Backurs, Arturs:** Talk Tue in Probabilistic inference 1, Poster Tue #50
- Bacry, Emmanuel:** Talk Wed in Causal Inference 1, Poster Wed #75
- Bagnell, Drew:** Talk Wed in Structured prediction, Poster Wed #121
- Bahdanau, Dzmitry:** Talk Tue in Recurrent neural networks 3, Poster Tue #77
- Bahri, Yasaman:** Talk Mon in Deep learning theory 1, Poster Tue #4
- Bailey, James:** Talk Tue in Recurrent neural networks 2, Poster Tue #31
- Balasubramanian, Krishnakumar:** Talk Wed in High dimensional estimation, Poster Wed #50
- Balcan, Nina:** Talk Mon in Transfer and multitask learning, Poster Mon #146, Talk Wed in Privacy and security 2, Poster Wed #86
- Balduzzi, David:** Talk Mon in Deep learning theory 1, Talk Tue in Game theory and multiagents, Poster Tue #124, Poster Tue #2, Poster Tue #5
- Ballas, Nicolas:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Balog, Matej:** Talk Tue in Probabilistic inference 3, Poster Tue #117
- Balzano, Laura:** Talk Mon in Active learning, Poster Mon #134, Talk Wed in Robust Estimation, Poster Wed #85
- Bamler, Robert:** Talk Wed in Language 1, Poster Wed #18
- Banerjee, Arindam:** Talk Wed in High dimensional estimation, Poster Wed #41, Poster Wed #58
- Baram, Nir:** Talk Mon in Reinforcement learning 3, Poster Mon #106, Talk Wed in Deep reinforcement learning 1, Poster Wed #99
- Barber, David:** Talk Mon in Continuous optimization 4, Poster Tue #22
- Bardenet, Rémi:** Talk Tue in Probabilistic inference 2, Poster Tue #80
- Bareinboim, Elias:** Talk Wed in Causal Inference 1, Talk Wed in Causal Inference 2, Poster Wed #93, Poster Wed #136
- Barreto, Andre:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Barrett, David:** Talk Tue in Deep learning 7: analysis, Poster Tue #113
- Bartlett, Peter:** Talk Mon in Deep learning theory 2, Poster Mon #47
- Barzilay, Regina:** Talk Mon in Deep learning 2: invariances, Poster Mon #67
- Bauer, Alexander:** Talk Wed in Privacy and security 1, Poster Wed #59
- Beckham, Christopher:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #88
- Begon, Jean-Michel:** Talk Tue in Ensemble methods, Poster Tue #86
- Belanger, David:** Talk Wed in Structured prediction, Poster Wed #130
- Belilovsky, Eugene:** Talk Mon in Probabilistic learning 3, Poster Mon #98
- Bellemare, Marc:** Talk Mon in Semisupervised and curriculum learning, Talk Mon in Reinforcement learning 4, Poster Mon #127, Poster Tue #23, Poster Tue #13, Talk Wed in Deep reinforcement learning 1, Poster Wed #64
- Bello, Irwan:** Talk Wed in Deep reinforcement learning 2, Poster Wed #134
- Below, Dan:** Talk Mon in Deep generative models 1, Poster Mon #10
- Ben Taieb, Souhaib:** Talk Wed in Time series, Poster Wed #87
- Bengio, Emmanuel:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Bengio, Yoshua:** Talk Mon in Deep learning theory 1, Talk Tue in Deep learning 7: analysis, Poster Tue #3, Poster Tue #105, Workshop Fri in C4.10
- Bengio, Samy:** Talk Mon in Deep learning theory 1, Poster Tue #3, Talk Wed in Deep learning 8: hardware, Poster Wed #17, Workshop Fri in C4.10
- Bennani, Younés:** Talk Tue in Clustering 2, Poster Wed #6
- Berg-Kirkpatrick, Taylor:** Talk Wed in Language 3, Poster Wed #123
- Bergmann, Urs:** Talk Mon in Deep generative models 1, Poster Mon #34
- Bernstein, Garrett:** Talk Wed in Privacy and security 1, Poster Wed #51, Workshop Fri in C4.5
- Beygelzimer, Alina:** Talk Mon in Online learning 4, Poster Tue #9
- Bhattacharya, Chiranjib:** Talk Tue in Clustering 2, Poster Wed #10
- Bhupatiraju, Surya:** Talk Tue in ML and programming, Poster Tue #118
- Bian, Jiang:** Talk Mon in Supervised learning 2, Poster Mon #48
- Bian, Andrew An:** Talk Wed in Combinatorial optimization 1, Poster Wed #72
- Bianchi, Matt:** Talk Wed in Healthcare, Poster Wed #60
- Bilge, Arman:** Talk Wed in Monte Carlo methods 2, Poster Wed #74
- Bishop, Adrian:** Talk Mon in Bayesian Nonparametrics, Poster Mon #141
- Blaschko, Matthew:** Talk Mon in Probabilistic learning 3, Poster Mon #98
- Blei, David:** Talk Mon in Probabilistic learning 1, Talk Mon in Deep generative models 3, Poster Mon #5, Poster Mon #21, Poster Mon #113, Workshop Thu in Parkside 1
- Blondel, Mathieu:** Talk Wed in Time series, Poster Wed #96
- Blundell, Charles:** Talk Mon in Reinforcement learning 3, Poster Mon #123, Talk Wed in Deep reinforcement learning 2, Poster Wed #125
- Blunsom, Phil:** Workshop Thu in C4.11
- Blunsom, Phil:** Talk Wed in Language 2, Poster Wed #62, Poster Wed #71
- Bogunovic, Ilija:** Talk Wed in Combinatorial optimization 1, Poster Wed #81
- Bojanowski, Piotr:** Talk Mon in Deep learning 2: invariances, Poster Mon #76, Talk Tue in Deep learning 6, Poster Wed #9
- Bolukbasi, Tolga:** Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #54
- Bonilla, Edwin:** Talk Wed in Gaussian processes, Poster Wed #126
- Boots, Byron:** Talk Wed in Continuous control, Talk Wed in Structured prediction, Poster Wed #121, Poster Wed #38
- Bora, Ashish:** Talk Tue in Sparsity 1, Poster Tue #103
- Botev, Aleksandar:** Talk Mon in Continuous optimization 4, Poster Tue #22
- Bottou, Léon:** Talk Mon in Deep generative models 2, Poster Mon #77
- Botvinick, Matthew:** Talk Mon in Reinforcement learning 3, Talk Mon in Deep learning 4: learning to learn, Poster Mon #123, Talk Tue in Deep learning 7: analysis, Poster Tue #113, Poster Tue #6
- Boubekki, Ahcène:** Talk Mon in Matrix factorization 1, Poster Mon #82
- Bourdev, Lubomir:** Talk Wed in Applications, Poster Wed #122
- Bowling, Michael:** Talk Mon in Reinforcement learning 4, Poster Tue #23
- Bošnjak, Matko:** Talk Tue in ML and programming, Poster Tue #126
- Brand, Daniel:** Talk Wed in Deep learning 8: hardware, Poster Wed #35
- Braut, Vincent:** Talk Tue in Clustering 2, Poster Wed #6
- Braun, Gábor:** Talk Tue in Continuous optimization 5, Poster Tue #32
- Braverman, Vladimir:** Talk Tue in Clustering 2, Poster Wed #12
- Brefeld, Ulf:** Talk Mon in Matrix factorization 1, Poster Mon #82
- Briol, Francois-Xavier:** Talk Tue in Probabilistic inference 2, Poster Tue #94
- Brockschmidt, Marc:** Talk Tue in ML and programming, Poster Tue #134
- Brown, Noam:** Talk Tue in Game theory and multiagents, Poster Tue #116
- Brunel, Victor-Emmanuel:** Talk Mon in Probabilistic learning 2, Poster Mon #62
- Brutzkus, Alon:** Talk Mon in Deep learning theory 3, Poster Mon #92
- Bubeck, Sebastian:** Talk Wed in Distributed optimization, Poster Wed #37
- Budden, David:** Talk Wed in Deep learning 8: hardware, Poster Wed #26
- Buhmann, Joachim:** Talk Wed in Combinatorial optimization 1, Poster Wed #72
- Bui, Hung:** Talk Mon in Deep generative models 3, Poster Mon #95, Talk Tue in Clustering 1, Poster Tue #69
- Bun, Mark:** Talk Wed in Privacy and security 1, Poster Wed #42
- Buntine, Wray:** Talk Tue in Networks and relational learning, Poster Tue #35, Workshop Fri in C4.11
- Burgess, Christopher:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Busa-Fekete, Robert:** Talk Mon in Online learning 1, Poster Mon #6
- Butler-Yeoman, Tony:** Talk Mon in Deep learning theory 1, Poster Tue #2
- Bärmann, Andreas:** Talk Mon in Online learning 4, Poster Tue #24
- C. Machado, Marlos:** Talk Mon in Reinforcement learning 4, Poster Tue #23
- Cai, Bryan:** Talk Wed in Privacy and security 1, Poster Wed #33, Workshop Fri in C4.5
- Cai, Deng:** Talk Tue in Sparsity 2, Poster Tue #111



- Calandriello, Daniele:** Talk Mon in Online learning 2, Poster Mon #81
- Campbell, Murray:** Talk Mon in Probabilistic learning 3, Poster Mon #107
- Canny, John:** Workshop Thu in C4.3
- Cao, Kris:** Workshop Thu in C4.11
- Carin, Lawrence:** Talk Tue in Networks and relational learning, Poster Tue #67, Talk Wed in Monte Carlo methods 2, Talk Wed in Language 2, Poster Wed #83, Poster Wed #89
- Carlson, David:** Talk Wed in Monte Carlo methods 1, Poster Wed #30
- Carmon, Daniel:** Talk Mon in Supervised learning 2, Poster Mon #57
- Carmon, Yair:** Talk Tue in Continuous optimization 7, Poster Tue #123
- Carr, Ambrose:** Workshop Thu in C4.4
- Carr, Peter:** Talk Tue in Game theory and multiagents, Poster Tue #132
- Carrière, Mathieu:** Talk Wed in Kernel methods, Poster Wed #120
- Caspi, Itai:** Talk Mon in Reinforcement learning 3, Poster Mon #106
- Castaldi, Peter:** Talk Tue in Clustering 2, Poster Wed #8
- Celis, Elisa:** Talk Tue in Sparsity 2, Poster Tue #143
- Cevher, Volkan:** Talk Wed in Combinatorial optimization 1, Poster Wed #81
- Cha, Moonsu:** Talk Mon in Deep generative models 2, Poster Mon #68
- Chakrabarti, Deepayan:** Talk Mon in Matrix factorization 1, Poster Mon #64
- Chan, Hubert:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #93
- Chandar, Sarath:** Workshop Thu in C4.1
- Chang, Sung-En:** Talk Mon in Latent feature models, Poster Mon #23
- Chang, Yale:** Talk Tue in Clustering 2, Poster Wed #8
- Chang, Kai-Wei:** Workshop Fri in C4.5
- Chanthirasegaran, pankajan:** Talk Wed in Language 2, Poster Wed #97
- Chau, Polo:** Workshop Thu in C4.3
- Chaudhry, Aditya:** Talk Wed in Causal Inference 1, Poster Wed #66
- Chaudhuri, Kamalika:** Talk Mon in Active learning, Poster Mon #138, Workshop Fri in C4.6
- Chebotar, Yevgen:** Talk Wed in Continuous control, Poster Wed #29
- Chechik, Gal:** Workshop Fri in C4.5
- Chen, Sheng:** Talk Wed in High dimensional estimation, Poster Wed #58
- Chen, Xi:** Talk Mon in Online learning 3, Poster Mon #125
- Chen, Junxiang:** Talk Tue in Clustering 2, Poster Wed #8
- Chen, Jianshu:** Talk Tue in Reinforcement learning 5, Poster Tue #41
- Chen, Jiecao:** Talk Mon in Online learning 3, Poster Mon #125
- Chen, Xixian:** Talk Wed in Robust Estimation, Poster Wed #67
- Chen, Changyou:** Talk Wed in Monte Carlo methods 2, Poster Wed #83
- Chen, Wilson Ye:** Talk Tue in Probabilistic inference 2, Poster Tue #94
- Chen, Bo:** Talk Wed in Monte Carlo methods 2, Poster Wed #100
- Chen, Guangyong:** Talk Tue in Probabilistic inference 3, Poster Tue #125
- Chen, Zhi:** Talk Wed in Language 2, Poster Wed #89
- Chen, Zhehui:** Talk Mon in Latent feature models, Poster Mon #31
- Chen, Jiecao:** Talk Wed in Kernel methods, Poster Wed #129
- Chen, Chao:** Talk Mon in Probabilistic learning 3, Poster Mon #116
- Chen, Yutian:** Talk Mon in Deep generative models 1, Talk Mon in Deep learning 4: learning to learn, Poster Mon #10, Poster Tue #6
- Chen, Bryant:** Talk Wed in Causal Inference 1, Poster Wed #93
- Chen, Wei:** Talk Mon in Supervised learning 2, Poster Mon #48, Talk Wed in Distributed optimization, Poster Wed #19
- Chen, Bryant:** Workshop Fri in C4.5
- Chen, Yichen:** Talk Mon in Continuous optimization 1, Poster Mon #27
- Chen, Bangrui:** Talk Mon in Online learning 3, Poster Mon #90
- Cherapanamjeri, Yeshwanth:** Talk Mon in Matrix factorization 2, Poster Mon #91
- Chierichetti, Flavio:** Talk Mon in Matrix factorization 3, Poster Tue #29
- Chintala, Soumith:** Talk Mon in Deep generative models 2, Poster Mon #77
- Cho, Michael:** Talk Tue in Clustering 2, Poster Wed #8
- Cho, Minsik:** Talk Wed in Deep learning 8: hardware, Poster Wed #35
- Choi, Arthur:** Talk Mon in Probabilistic learning 3, Poster Mon #124
- Choi, Seungjin:** Talk Tue in Networks and relational learning, Poster Tue #43
- Choromanska, Anna:** Talk Mon in Supervised learning 1, Poster Mon #16
- Chorowski, Jan:** Talk Tue in Recurrent neural networks 2, Poster Tue #63
- Chou, Po-Wei:** Talk Wed in Continuous control, Poster Wed #20
- Chrzanowski, Mike:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Cisse, Moustapha:** Talk Wed in Deep learning 8: hardware, Poster Wed #53
- Cisse, Moustapha:** Talk Tue in Deep learning 6, Poster Wed #9
- Coates, Adam:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Cockayne, Jon:** Talk Tue in Probabilistic inference 2, Poster Tue #94
- Cong, Yulai:** Talk Wed in Monte Carlo methods 2, Poster Wed #100
- Converse, Geoff:** Talk Mon in Active learning, Poster Mon #130
- Cortes, Corinna:** Talk Mon in Deep learning 3: metalearning, Poster Mon #121, Workshop Fri in C4.6
- Courville, Aaron:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Couso, Ines:** Talk Tue in Ranking and preferences, Poster Tue #81
- Cutajar, Kurt:** Talk Wed in Gaussian processes, Poster Wed #126
- Cuturi, Marco:** Talk Wed in Time series, Talk Wed in Kernel methods, Poster Wed #120, Poster Wed #96
- Czarnecki, Wojciech:** Talk Mon in Deep learning 1: backprop, Poster Mon #1, Poster Mon #9
- D'Eramo, Carlo:** Talk Mon in Reinforcement learning 1, Poster Mon #28
- Dabney, Will:** Talk Mon in Reinforcement learning 4, Poster Tue #13
- Dahl, George:** Talk Wed in Applications, Poster Wed #131
- Dai, Hajun:** Talk Tue in Networks and relational learning, Poster Tue #59
- Dai, Zhenwen:** Talk Wed in Bayesian Optimization, Poster Wed #22
- Dai, Bo:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #120, Talk Wed in Large scale learning, Poster Wed #119
- Damianou, Andreas:** Talk Wed in Bayesian Optimization, Poster Wed #22
- Danihelka, Ivo:** Talk Mon in Deep generative models 1, Poster Mon #18
- Darrell, Trevor:** Talk Mon in Reinforcement learning 3, Poster Mon #88
- Darwiche, Adnan:** Talk Mon in Probabilistic learning 3, Poster Mon #124
- Das, HIRAKENDU:** Talk Tue in Information theory, Poster Tue #128
- Dasarathy, Gautam:** Talk Wed in Bayesian Optimization, Poster Wed #49
- Dasgupta, Sanjoy:** Talk Mon in Active learning, Poster Mon #145
- Daskalakis, Constantinos:** Talk Wed in Privacy and security 1, Poster Wed #33
- Daumé, Hal:** Talk Mon in Supervised learning 2, Poster Mon #75
- Daumé III, Hal:** Talk Mon in Active learning, Poster Mon #142
- Dauphin, Yann:** Talk Tue in Deep learning 6, Talk Wed in Language 3, Poster Wed #114, Poster Wed #9, Poster Wed #105
- Davidson, James:** Talk Mon in Reinforcement learning 1, Poster Mon #4
- Davies, Michael:** Talk Mon in Continuous optimization 2, Poster Mon #69
- Dawson, Colin:** Talk Mon in Bayesian Nonparametrics, Poster Mon #129
- Daxberger, Erik:** Talk Wed in Gaussian processes, Poster Wed #108
- de Balle Pigem, Borja:** Talk Tue in Spectral methods, Poster Tue #52
- De Castro, Yohann:** Talk Mon in Matrix factorization 1, Poster Mon #73
- de Freitas, Nando:** Talk Mon in Deep generative models 1, Talk Mon in Deep learning 4: learning to learn, Poster Mon #10, Poster Tue #6, Poster Tue #11
- Dean, Jeff:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Degrís, Thomas:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Dekel, Ofer:** Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #54
- Delbruck, Tobi:** Talk Mon in Recurrent neural networks 1, Poster Mon #136
- Dembczynski, Krzysztof:** Talk Tue in Learning theory 2, Poster Tue #104
- Dempsey, Walter:** Talk Wed in Healthcare, Poster Wed #52
- Deng, Yuntian:** Talk Tue in Probabilistic inference 3, Poster Tue #141
- Deng, Li:** Talk Tue in Recurrent neural networks 4, Poster Tue #130
- Deng, Yuntian:** Talk Wed in Language 1, Poster Wed #45
- Denil, Misha:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #6, Poster Tue #11
- Dennis, Michael:** Talk Wed in Healthcare, Poster Wed #52
- DeSalvo, Giulia:** Workshop Fri in C4.6
- Devlin, Jacob:** Talk Tue in ML and programming, Poster Tue #118
- Dey, Debadepta:** Talk Mon in Online learning 3, Poster Mon #117
- Dhillon, Inderjit:** Talk Mon in Deep learning theory 2, Talk Mon in Continuous optimization 3, Poster Mon #47, Poster Mon #105, Talk Tue in Ensemble methods, Poster Tue #79
- Diakonikolas, Ilias:** Talk Wed in High dimensional estimation, Poster Wed #23
- Diamos, Gregory:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Dick, Travis:** Talk Wed in Privacy and security 2, Poster Wed #86
- Dieleman, Sander:** Talk Tue in Recurrent neural networks 3, Poster Tue #98
- Dill, David:** Talk Tue in ML and programming, Poster Tue #142
- Dimakis, Alexandros:** Talk Tue in Probabilistic inference 1, Talk Tue in Sparsity 1, Talk Tue in Information theory, Poster Tue #136, Poster Tue #34, Poster Tue #103, Talk Wed in Causal Inference 1, Talk Wed in Combinatorial optimization 1, Talk Wed in Causal Inference 2,
- Poster Wed #98, Poster Wed #84, Poster Wed #118
- Dinh, Vu:** Talk Wed in Monte Carlo methods 2, Poster Wed #74
- Dinh, Laurent:** Talk Mon in Deep learning theory 1, Poster Tue #3
- Djatkiko, Mentari:** Workshop Fri in C4.5
- Domingos, Pedro:** Workshop Thu in C4.5
- Domke, Justin:** Talk Tue in Probabilistic inference 2, Poster Tue #87
- Donahue, Christopher:** Talk Wed in Applications, Poster Wed #104
- Dong, Xuanyi:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #102
- Donini, Michele:** Talk Tue in Continuous optimization 6, Poster Tue #92
- Donnelly, Peter:** Invited Talk Tue in Darling Harbour Theater
- Dorfman, Nimrod:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Doshi-Velez, Finale:** Tutorial Sun in Tutorials Session B
- Du, Lan:** Talk Tue in Networks and relational learning, Poster Tue #35
- Du, Simon:** Talk Tue in Reinforcement learning 5, Poster Tue #41
- du Plessis, Marthinus:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #111
- Dubcek, Tena:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Dubhashi, Devdatt:** Talk Tue in Clustering 2, Poster Wed #10
- Duchi, John:** Talk Tue in Continuous optimization 6, Talk Tue in Continuous optimization 7, Poster Tue #99, Poster Tue #123
- Dudík, Miroslav:** Talk Tue in Reinforcement learning 5, Poster Tue #49
- Dukkkipati, Ambedkar:** Talk Mon in Recurrent neural networks 1, Poster Mon #128
- Dulac-Arnold, Gabriel:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Dy, Jennifer:** Talk Tue in Clustering 2, Poster Wed #8
- Dyer, Chris:** Workshop Thu in C4.11
- E, Weinan:** Talk Tue in Continuous optimization 6, Poster Tue #71
- Earle, Adam:** Talk Mon in Reinforcement learning 4, Poster Tue #18
- Eck, Douglas:** Talk Tue in Recurrent neural networks 3, Poster Tue #77, Poster Tue #70, Poster Tue #98
- Eckstein, Jonathan:** Talk Wed in Combinatorial optimization 2, Poster Wed #141
- Efros, Alexei:** Talk Mon in Reinforcement learning 3, Poster Mon #88
- Eldawy, Mohamed:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Elenberg, Ethan:** Talk Wed in Combinatorial optimization 1, Poster Wed #98



- Elmahdy, Adel:** Talk Tue in Ranking and preferences, Poster Tue #102
- Engel, Jesse:** Talk Tue in Recurrent neural networks 3, Poster Tue #98
- Ermon, Stefano:** Talk Mon in Deep generative models 3, Poster Mon #86
- Etrue Howard, Evans:** Talk Mon in Online learning 3, Poster Mon #99
- Fadnis, Kshitij:** Talk Mon in Probabilistic learning 3, Poster Mon #107
- Falahatgar, Moein:** Talk Tue in Ranking and preferences, Poster Tue #95
- Fan, Linxi:** Talk Wed in Applications, Poster Wed #113
- Fan, Xiangmin:** Workshop Thu in C4.3
- Fan, Kai:** Talk Wed in Language 2, Poster Wed #89
- Fan, Angela:** Talk Wed in Language 3, Poster Wed #105
- Farajtabar, Mehrdad:** Talk Mon in Reinforcement learning 2, Poster Mon #79
- Farina, Gabriele:** Talk Tue in Game theory and multiagents, Poster Tue #108
- Farquhar, Gregory:** Talk Wed in Deep reinforcement learning 1, Poster Wed #82
- Feldman, Dan:** Talk Mon in Matrix factorization 1, Poster Mon #46
- Figureiredo, Mario:** Talk Wed in Distributed optimization, Poster Wed #28
- Filippone, Maurizio:** Talk Wed in Gaussian processes, Poster Wed #126
- Finn, Chelsea:** Tutorial Sun in Tutorials Session B, Talk Mon in Deep learning 3: metalearning, Poster Mon #112
- Fischer, Asja:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Foerster, Jakob:** Talk Tue in Recurrent neural networks 2, Poster Tue #63, Talk Wed in Deep reinforcement learning 1, Poster Wed #82
- Forney, Andrew:** Talk Wed in Causal Inference 2, Poster Wed #136
- Foti, Nicholas:** Talk Tue in Probabilistic inference 3, Poster Tue #109, Talk Wed in Monte Carlo methods 2, Poster Wed #92
- Fountoulakis, Kimon:** Talk Tue in Spectral methods, Poster Tue #60
- Fox, Emily:** Talk Wed in Monte Carlo methods 2, Poster Wed #92
- Frahling, Gereon:** Talk Tue in Clustering 2, Poster Wed #12
- Franceschi, Luca:** Talk Tue in Continuous optimization 6, Poster Tue #92
- Frasconi, Paolo:** Talk Tue in Continuous optimization 6, Poster Tue #92
- Frazier, Peter:** Talk Mon in Online learning 3, Poster Mon #90
- Frean, Marcus:** Talk Mon in Deep learning theory 1, Poster Tue #5
- Frossard, Pascal:** Talk Mon in Deep learning 2: invariances, Poster Mon #58
- Fu, Xiao:** Talk Tue in Clustering 1, Poster Tue #53
- Fujimaki, Ryohei:** Talk Tue in Sparsity 1, Poster Tue #89
- Fukuchi, Kazuto:** Talk Wed in Privacy and security 2, Poster Wed #95
- Futoma, Joseph:** Talk Wed in Healthcare, Poster Wed #43
- Gaertner, Thomas:** Talk Wed in Kernel methods, Poster Wed #143
- Gal, Yarin:** Talk Mon in Probabilistic learning 2, Poster Mon #71, Talk Wed in Deep learning 9: probabilistic, Poster Wed #61
- Gan, Zhe:** Talk Wed in Monte Carlo methods 2, Talk Wed in Language 2, Poster Wed #83, Poster Wed #89
- Ganguli, Surya:** Talk Mon in Deep learning theory 3, Poster Mon #110, Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #46
- Gao, Tian:** Talk Mon in Probabilistic learning 3, Poster Mon #107
- Gao, Wenbo:** Talk Mon in Continuous optimization 4, Poster Tue #17
- Garber, Dan:** Talk Mon in Latent feature models, Poster Mon #7
- Garnett, Roman:** Talk Mon in Active learning, Poster Mon #130, Workshop Thu in C4.9
- Gaunt, Alex:** Talk Tue in ML and programming, Poster Tue #134
- Gautier, Guillaume:** Talk Tue in Probabilistic inference 2, Poster Tue #80
- Gaïffas, Stéphane:** Talk Wed in Causal Inference 1, Poster Wed #75
- Ge, Rong:** Talk Mon in Deep generative models 2, Talk Mon in Matrix factorization 2, Poster Mon #109, Poster Mon #41, Talk Tue in Continuous optimization 7, Poster Tue #139
- Ge, Dongdong:** Talk Mon in Continuous optimization 1, Poster Mon #27
- Gehring, Jonas:** Talk Wed in Language 3, Poster Wed #114
- Geiger, Andreas:** Talk Mon in Deep generative models 1, Poster Mon #26
- Gens, Robert:** Workshop Thu in C4.5
- Gentile, Claudio:** Talk Mon in Online learning 3, Poster Mon #99
- George, Dileep:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Gerz, Daniela:** Workshop Thu in C4.11
- Geurts, Pierre:** Talk Tue in Ensemble methods, Poster Tue #86
- Ghahramani, Zoubin:** Talk Mon in Probabilistic learning 1, Talk Mon in Probabilistic learning 2, Talk Mon in Bayesian Nonparametrics, Poster Mon #71, Poster Mon #29, Poster Mon #144, Talk Tue in Networks and relational learning, Talk Tue in Probabilistic inference 3, Poster Tue #117, Poster Tue #43, Talk Wed in Monte Carlo methods 2, Poster Wed #65
- Ghavamzadeh, Mohammad:** Talk Mon in Online learning 1, Talk Mon in Deep generative models 3, Talk Mon in Online learning 4, Poster Mon #95, Poster Mon #22, Poster Mon #30, Poster Tue #14
- Gheslghi Azar, Mohammad:** Talk Mon in Reinforcement learning 1, Poster Mon #12
- Ghosh, Joydeep:** Talk Wed in Combinatorial optimization 1, Poster Wed #98
- Gibiasky, Andrew:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Gifford, David:** Talk Tue in Recurrent neural networks 4, Poster Tue #114
- Gilboa, Dar:** Talk Wed in Monte Carlo methods 1, Poster Wed #30
- Gilmer, Justin:** Talk Tue in Recurrent neural networks 2, Poster Tue #63, Talk Wed in Applications, Poster Wed #131
- Girolami, Mark:** Talk Tue in Probabilistic inference 2, Poster Tue #94
- Gittens, Alex:** Talk Mon in Continuous optimization 2, Poster Mon #51
- Gittens, Alex:** Talk Tue in Learning theory 1, Poster Wed #1
- Globerson, Amir:** Talk Mon in Supervised learning 2, Talk Mon in Deep learning theory 3, Poster Mon #92, Poster Mon #57
- Glowacka, Dorota:** Workshop Fri in C4.11
- Goel, Vaibhava:** Talk Mon in Deep generative models 2, Poster Mon #50
- GOH, CHI:** Talk Mon in Transfer and multitask learning, Poster Mon #135
- Golbabaee, Mohammad:** Talk Mon in Continuous optimization 2, Poster Mon #69
- Goldberg, Noam:** Talk Wed in Combinatorial optimization 2, Poster Wed #141
- Goldfarb, Donald:** Talk Mon in Continuous optimization 4, Poster Tue #17
- Goldstein, Tom:** Talk Mon in Continuous optimization 3, Poster Mon #122, Talk Wed in Distributed optimization, Poster Wed #28
- Gollapudi, Sreenivas:** Talk Mon in Matrix factorization 3, Poster Tue #29
- Gonzalvo, Xavi:** Talk Mon in Deep learning 3: metalearning, Poster Mon #121
- González, Javier:** Talk Wed in Bayesian Optimization, Poster Wed #22, Poster Wed #40
- Goodfellow, Ian:** Workshop Thu in Parkside 1
- Gopalan, Aditya:** Talk Mon in Online learning 2, Poster Mon #72
- Gordon, Geoff:** Talk Wed in Structured prediction, Poster Wed #121
- Gorham, Jackson:** Talk Tue in Probabilistic inference 2, Poster Tue #101
- Goude, Yannig:** Talk Mon in Matrix factorization 1, Poster Mon #73
- Goyal, Ankit:** Talk Tue in Metric learning, Poster Wed #16
- Goyal, Anirudh:** Workshop Fri in C4.10
- Goyal, Saurabh:** Talk Mon in Supervised learning 1, Poster Mon #24, Talk Tue in Metric learning, Poster Wed #16
- Grangier, David:** Talk Wed in Deep learning 8: hardware, Talk Wed in Language 3, Poster Wed #53, Poster Wed #114, Poster Wed #105
- Grave, Edouard:** Talk Tue in Deep learning 6, Poster Wed #9
- Grave, Edouard:** Talk Wed in Deep learning 8: hardware, Poster Wed #53
- Graves, Alex:** Talk Mon in Deep learning 1: backprop, Talk Mon in Deep generative models 1, Talk Mon in Semisupervised and curriculum learning, Poster Mon #1, Poster Mon #127, Poster Mon #18
- Gray, Andrew:** Tutorial Sun in Tutorials Session C, Workshop Thu in C4.10
- Greenside, Peyton:** Talk Mon in Deep learning 1: backprop, Poster Mon #25
- Grefenstette, Edward:** Talk Wed in Language 2, Poster Wed #71
- Gretton, Arthur:** Talk Wed in Kernel methods, Poster Wed #111
- Grossglauer, Matthias:** Talk Tue in Ranking and preferences, Poster Tue #74, Poster Tue #88
- Gu, Ming:** Talk Mon in Matrix factorization 3, Poster Tue #25
- Gu, Quanquan:** Talk Mon in Matrix factorization 3, Poster Tue #20, Talk Wed in Causal Inference 1, Talk Wed in Robust Estimation, Poster Wed #94, Poster Wed #76, Poster Wed #66
- Gu, Shixiang:** Talk Tue in Recurrent neural networks 3, Poster Tue #77
- Guestin, Carlos:** Talk Tue in Continuous optimization 5, Poster Tue #64
- Guez, Arthur:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Gui, Shupeng:** Talk Tue in Sparsity 1, Poster Tue #89
- Guo, Ruiqi:** Talk Wed in Large scale learning, Poster Wed #119
- Guo, Chuan:** Talk Tue in Deep learning 7: analysis, Poster Tue #137
- Gupta, Abhinav:** Talk Mon in Reinforcement learning 1, Poster Mon #4
- Gupta, Shubham:** Talk Mon in Recurrent neural networks 1, Poster Mon #128
- Gupta, Kartik:** Talk Mon in Matrix factorization 2, Poster Mon #91
- Gupta, Sunil:** Talk Wed in Gaussian processes, Poster Wed #142
- Gupta, Chirag:** Talk Tue in Metric learning, Poster Wed #16
- Gustafson, David:** Talk Wed in Healthcare, Poster Wed #52
- Gygli, Michael:** Talk Wed in Structured prediction, Poster Wed #112
- Gómez Colmenarejo, Sergio:** Talk Mon in Deep generative models 1, Talk Mon in Deep learning 4: learning to learn, Poster Mon #10, Poster Tue #6, Poster Tue #11
- Görmitz, Nico:** Talk Wed in Privacy and security 1, Poster Wed #59
- Haarnoja, Tuomas:** Talk Mon in Reinforcement learning 2, Poster Mon #52
- Hadfield-Menell, Dylan:** Workshop Fri in C4.7
- HADJERES, Gaëtan:** Talk Tue in Recurrent neural networks 3, Poster Tue #91
- Hadsell, Raia:** Invited Talk Wed in Darling Harbour Theater
- Hajinezhad, Davood:** Talk Mon in Continuous optimization 3, Poster Mon #87
- Hallak, Assaf:** Talk Tue in Reinforcement learning 5, Poster Tue #57
- Han, Insu:** Talk Tue in Probabilistic inference 2, Poster Tue #73
- Hanna, Josiah:** Talk Tue in Reinforcement learning 5, Poster Tue #33
- Harada, Tatsuya:** Talk Mon in Transfer and multitask learning, Poster Mon #131
- Harandi, Mehrtaash:** Talk Tue in Metric learning, Poster Wed #15
- Hariharan, Sanjay:** Talk Wed in Healthcare, Poster Wed #43
- Harley, Tim:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Hartford, Jason:** Talk Wed in Causal Inference 2, Poster Wed #127
- Hartley, Richard:** Talk Tue in Metric learning, Poster Wed #15
- Hassabis, Demis:** Talk Wed in Deep reinforcement learning 2, Poster Wed #125
- Hassani, Hamed:** Talk Tue in Learning theory 2, Poster Tue #76
- Hassidim, Avinatan:** Talk Wed in Combinatorial optimization 1, Poster Wed #63
- Hausman, Karol:** Talk Wed in Continuous control, Poster Wed #29
- Hay, Michael:** Talk Wed in Privacy and security 1, Poster Wed #51
- Hayashi, Kohei:** Talk Mon in Matrix factorization 2, Poster Mon #126
- Hazan, Elad:** Talk Mon in Online learning 2, Poster Mon #45
- He, Lifang:** Talk Mon in Supervised learning 1, Poster Mon #39
- He, Niao:** Talk Wed in Large scale learning, Poster Wed #119
- He, Xiaofoe:** Talk Tue in Sparsity 2, Poster Tue #111
- He, Bryan:** Talk Mon in Probabilistic learning 3, Poster Mon #89
- Heaukulani, Creighton:** Talk Tue in Networks and relational learning, Poster Tue #43
- Heckel, Reinhard:** Talk Mon in Online learning 1, Poster Mon #38
- Heess, Nicolas:** Talk Wed in Deep reinforcement learning 2, Poster Wed #107
- Hein, Matthias:** Talk Mon in Deep learning theory 1, Talk Mon in Online learning 4, Poster Tue #1, Poster Tue #28
- Held, David:** Talk Mon in Reinforcement learning 2, Poster Mon #43



- Heller, Katherine:** Talk Wed in Healthcare, Poster Wed #43
- Hellicar, Andrew:** Talk Tue in Recurrent neural networks 2, Poster Tue #31
- Henao, Ricardo:** Talk Wed in Monte Carlo methods 2, Talk Wed in Language 2, Poster Wed #83, Poster Wed #89
- Heng, Pheng Ann:** Talk Tue in Probabilistic inference 3, Poster Tue #125
- Henriques, Joao:** Talk Mon in Deep learning 2: invariances, Poster Mon #49
- Henzinger, Monika:** Talk Tue in Spectral methods, Poster Tue #60
- Hernandez, Jonathan:** Talk Wed in Applications, Poster Wed #113
- Hernandez-Lobato, Daniel:** Talk Wed in Gaussian processes, Poster Wed #117
- Hernandez-Lobato, Jose:** Talk Tue in Recurrent neural networks 3, Poster Tue #77, Talk Wed in Bayesian Optimization, Talk Wed in Language 3, Poster Wed #57, Poster Wed #132
- Hershey, John:** Talk Wed in Language 1, Poster Wed #54
- Hessel, Matteo:** Talk Tue in Deep reinforcement learning 1, Poster Wed #91
- Higgins, Irina:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Hinder, Oliver:** Talk Tue in Continuous optimization 7, Poster Tue #123
- Hirayama, Jun-ichiro:** Talk Mon in Latent feature models, Poster Mon #15
- Hirche, Sandra:** Talk Wed in Continuous control, Poster Wed #47
- Ho, Nhat:** Talk Tue in Clustering 1, Poster Tue #69
- Ho, Mark:** Talk Mon in Reinforcement learning 3, Poster Mon #97
- Hoffman, Matthew:** Talk Mon in Deep generative models 3, Poster Mon #104
- Hoffman, Matthew:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #6, Poster Tue #11
- Hoi, Steven:** Talk Wed in Distributed optimization, Poster Wed #46
- Holmes, Christopher:** Talk Mon in Probabilistic learning 2, Poster Mon #80
- Hong, Bin:** Talk Tue in Sparsity 2, Poster Tue #111
- Hong, Mingyi:** Talk Mon in Continuous optimization 3, Poster Mon #87, Talk Tue in Clustering 1, Poster Tue #53
- Honkela, Antti:** Workshop Fri in C4.4, Workshop Fri in C4.5
- Hori, Takaaki:** Talk Wed in Language 1, Poster Wed #54
- Hornakova, Andrea:** Talk Wed in Combinatorial optimization 2, Poster Wed #115
- How, Jonathan:** Talk Tue in Game theory and multiagents, Poster Tue #140
- Hsieh, Cho-Jui:** Talk Tue in Ensemble methods, Poster Tue #79
- Hu, Shuguang:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #93
- Hu, Weihua:** Talk Tue in Infomation theory, Poster Tue #144
- Hu, Changwei:** Talk Tue in Networks and relational learning, Poster Tue #67
- Hu, Scott:** Talk Wed in Healthcare, Poster Wed #34
- Hu, Hao:** Talk Mon in Recurrent neural networks 1, Poster Mon #132
- Hu, Bin:** Talk Tue in Continuous optimization 6, Poster Tue #78
- Hu, Zhiting:** Talk Wed in Language 2, Talk Wed in Language 3, Poster Wed #80, Poster Wed #123
- Huang, Hui:** Talk Tue in Probabilistic inference 3, Poster Tue #125
- Huang, Chaofan:** Talk Mon in Bayesian Nonparametrics, Poster Mon #129
- Huang, Bert:** Workshop Fri in C4.5
- Huang, Po-Sen:** Talk Tue in Recurrent neural networks 4, Poster Tue #130
- Huang, Tzu-Kuo:** Talk Mon in Active learning, Poster Mon #142
- Hughes, Michael:** Talk Mon in Bayesian Nonparametrics, Poster Mon #137
- Humayun, Ahmad:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #120
- Huynh, Viet:** Talk Tue in Clustering 1, Poster Tue #69
- Hwang, Changho:** Talk Tue in Ensemble methods, Poster Tue #100
- Hwang, Sung Ju:** Talk Mon in Deep learning 3: metalearning, Poster Mon #103, Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #62
- Hyndman, Rob:** Talk Wed in Time series, Poster Wed #87
- Hyvärinen, Aapo:** Talk Mon in Latent feature models, Poster Mon #15
- Hébrail, Georges:** Talk Mon in Matrix factorization 1, Poster Mon #73
- Höner, János:** Talk Wed in Privacy and security 1, Poster Wed #59
- Hüllermeier, Eyke:** Talk Tue in Ranking and preferences, Poster Tue #81
- Ibrahim, Ahmed M. Alaa:** Talk Wed in Healthcare, Poster Wed #34
- Imaizumi, Masaaki:** Talk Mon in Matrix factorization 2, Poster Mon #126
- Ingraham, John:** Talk Tue in Probabilistic inference 1, Poster Tue #58
- Islam, Riashat:** Talk Mon in Probabilistic learning 2, Poster Mon #71
- Ithapu, Vamsi:** Talk Wed in Healthcare, Poster Wed #25
- Jaakkola, Tommi:** Talk Mon in Deep learning 2: invariances, Poster Mon #67, Talk Tue in Recurrent neural networks 4, Poster Tue #114, Talk Wed in Healthcare, Poster Wed #60
- Jabbari, Shahin:** Talk Mon in Reinforcement learning 1, Poster Mon #20
- Jaderberg, Max:** Talk Mon in Deep learning 1: backup, Poster Mon #1, Poster Mon #9, Talk Wed in Deep reinforcement learning 2, Poster Wed #107
- Jaggi, Martin:** Talk Tue in Continuous optimization 5, Poster Tue #56
- Jain, Vikas:** Talk Wed in Structured prediction, Poster Wed #144
- Jain, Prateek:** Talk Mon in Deep learning theory 2, Talk Mon in Matrix factorization 2, Talk Mon in Active learning, Poster Mon #91, Poster Mon #138, Poster Mon #47, Talk Tue in Metric learning, Poster Wed #16, Workshop Thu in C4.7
- Jaitly, Navdeep:** Tutorial Sun in Tutorials Session B
- Jalal, Ajil:** Talk Tue in Sparsity 1, Poster Tue #103
- James, Lancelot:** Talk Tue in Networks and relational learning, Poster Tue #43
- Jaques, Natasha:** Talk Tue in Recurrent neural networks 3, Poster Tue #77
- Jastrzebski, Stanislaw:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Jegelka, Stefanie:** Talk Mon in Bayesian Nonparametrics, Poster Mon #133, Talk Wed in Bayesian Optimization, Talk Wed in Combinatorial optimization 2, Poster Wed #31, Poster Wed #133
- Jegou, Herve:** Talk Wed in Deep learning 8: hardware, Poster Wed #53
- Jenatton, Rodolphe:** Talk Wed in Bayesian Optimization, Poster Wed #40
- Jernite, Yacine:** Talk Mon in Supervised learning 1, Poster Mon #16
- Jetchev, Nikolay:** Talk Mon in Deep generative models 1, Poster Mon #34
- Ji, Geng:** Talk Mon in Bayesian Nonparametrics, Poster Mon #137
- Jiang, Shunhua:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #16
- Jiang, Heinrich:** Talk Tue in Learning theory 2, Poster Tue #83, Poster Tue #90
- Jiang, Shali:** Talk Mon in Active learning, Poster Mon #130
- Jiang, Nan:** Talk Tue in Reinforcement learning 5, Poster Tue #65
- Jiang, Biye:** Workshop Thu in C4.3
- Jin, Chi:** Talk Mon in Matrix factorization 2, Poster Mon #109, Talk Tue in Continuous optimization 7, Poster Tue #139
- Jin, Wengong:** Talk Mon in Deep learning 2: invariances, Poster Mon #67
- Jin, Zhi:** Talk Wed in Language 1, Poster Wed #36
- Jing, Li:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Jitkrittum, Wittawat:** Talk Wed in Kernel methods, Poster Wed #111
- Johansson, Fredrik:** Talk Tue in Clustering 2, Talk Wed in Causal Inference 1, Poster Wed #10, Poster Wed #101
- Johnson, Sterling:** Talk Wed in Healthcare, Poster Wed #25
- Johnson, Tyler:** Talk Tue in Continuous optimization 5, Poster Tue #64
- Joly, Arnaud:** Talk Tue in Ensemble methods, Poster Tue #86
- Jordan, Michael:** Talk Mon in Continuous optimization 2, Talk Mon in Deep learning 3: metalearning, Poster Mon #85, Poster Mon #51, Talk Tue in Continuous optimization 7, Poster Tue #139
- Joseph, Matthew:** Talk Mon in Reinforcement learning 1, Poster Mon #20
- Joulin, Armand:** Talk Mon in Deep learning 2: invariances, Poster Mon #76, Talk Wed in Deep learning 8: hardware, Poster Wed #53
- Kadoury, Samuel:** Talk Tue in Recurrent neural networks 2, Poster Tue #39
- Kagawa, Ai:** Talk Wed in Combinatorial optimization 2, Poster Wed #141
- Kajino, Hiroshi:** Talk Wed in Time series, Poster Wed #69
- Kakade, Sham:** Talk Tue in Continuous optimization 7, Poster Tue #139, Workshop Thu in C4.5
- Kakizaki, Kazuya:** Talk Wed in Privacy and security 2, Poster Wed #95
- Kalchbrenner, Nal:** Talk Mon in Deep generative models 1, Poster Mon #10, Poster Mon #18
- Kale, Satyen:** Talk Mon in Online learning 4, Poster Tue #19
- Kallus, Nathan:** Talk Wed in Causal Inference 2, Poster Wed #109
- Kalousis, Alexandros:** Talk Tue in Deep learning 6, Poster Wed #11
- Kamath, Gautam:** Talk Wed in High dimensional estimation, Talk Wed in Privacy and security 1, Poster Wed #33, Poster Wed #23
- Kambadur, Prabhanjan:** Talk Tue in Probabilistic inference 2, Poster Tue #73
- kandasamy, kirthevasan:** Talk Wed in Bayesian Optimization, Poster Wed #49
- Kane, Daniel:** Talk Wed in High dimensional estimation, Poster Wed #23
- Kanervisto, Anssi:** Talk Wed in Language 1, Poster Wed #45
- Kang, Yongguo:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Kansky, Ken:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Kanwal, Maxinder:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Kapoor, Ashish:** Talk Mon in Online learning 3, Poster Mon #117
- Kapralov, Michael:** Talk Mon in Supervised learning 2, Poster Mon #66
- Kar, Purushottam:** Talk Mon in Online learning 3, Poster Mon #99
- Kara, Kaan:** Talk Wed in Large scale learning, Poster Wed #128
- Karampatziakis, Nikos:** Talk Mon in Supervised learning 2, Poster Mon #75, Talk Tue in Infomation theory, Poster Tue #136
- Karatzoglou, Alexandros:** Talk Mon in Online learning 3, Poster Mon #99
- Karbasi, Amin:** Talk Wed in Privacy and security 1, Talk Wed in Combinatorial optimization 1, Talk Wed in Combinatorial optimization 2, Poster Wed #42, Poster Wed #106, Poster Wed #90
- Karnin, Zohar:** Talk Mon in Online learning 4, Poster Tue #19
- Karpathy, Andrej:** Talk Wed in Applications, Poster Wed #113
- Kaski, Samuel:** Workshop Fri in C4.4
- Kastner, Kyle:** Talk Mon in Probabilistic learning 3, Poster Mon #98
- Katabi, Dina:** Talk Wed in Healthcare, Poster Wed #60
- Katt, Sammie:** Talk Mon in Reinforcement learning 3, Poster Mon #115
- Kavukcuoglu, Koray:** Talk Mon in Deep learning 1: backup, Talk Mon in Deep generative models 1, Talk Mon in Semisupervised and curriculum learning, Poster Mon #1, Poster Mon #9, Poster Mon #18, Poster Mon #127, Talk Wed in Deep reinforcement learning 2, Poster Wed #107
- Kawanabe, Motoaki:** Talk Mon in Latent feature models, Poster Mon #15
- Ke, Yiping:** Talk Mon in Transfer and multitask learning, Poster Mon #135
- Ke, Nan:** Workshop Fri in C4.10
- Ke, Chuyang:** Talk Tue in Sparsity 1, Poster Tue #89
- Kearns, Michael:** Talk Mon in Reinforcement learning 1, Poster Mon #20, Talk Tue in Learning theory 1, Poster Wed #3
- Keerthi, Sathya:** Talk Tue in Ensemble methods, Poster Tue #79
- Khalil, Elias:** Talk Mon in Reinforcement learning 2, Poster Mon #79
- Khanna, Aran:** Tutorial Sun in Tutorials Session A
- KHANNA, RAJIV:** Talk Wed in Combinatorial optimization 1, Poster Wed #98
- Khasanova, Renata:** Talk Mon in Deep learning 2: invariances, Poster Mon #58
- Kim, Yongjune:** Talk Mon in Deep learning theory 2, Poster Mon #65
- Kim, Gunhee:** Talk Mon in Deep learning 3: metalearning, Poster Mon #103
- Kim, Juyong:** Talk Mon in Deep learning 3: metalearning, Poster Mon #103
- Kim, Jiwon:** Talk Mon in Deep generative models 2, Poster Mon #68
- Kim, Been:** Tutorial Sun in Tutorials Session B, Workshop Thu in C4.8
- Kim, Taeksoo:** Talk Mon in Deep generative models 2, Poster Mon #68



- Kim, Hyunsoo:** Talk Mon in Deep generative models 2, Poster Mon #68
- King, Irwin:** Talk Wed in Robust Estimation, Poster Wed #67
- Klein, Dan:** Talk Mon in Reinforcement learning 4, Poster Tue #27
- Kleinberg, Jon:** Talk Mon in Deep learning theory 3, Poster Mon #110
- Klivans, Adam:** Talk Tue in Probabilistic inference 1, Poster Tue #34
- Knowles, David:** Talk Mon in Bayesian Nonparametrics, Poster Mon #144
- Kocaoglu, Murat:** Talk Wed in Causal Inference 1, Poster Wed #84
- Koh, Pang Wei:** Talk Mon in Supervised learning 2, Poster Mon #84
- Kohler, Jonas:** Talk Tue in Continuous optimization 7, Poster Tue #107
- Kohli, Pushmeet:** Talk Mon in Bayesian Nonparametrics, Poster Mon #133, Talk Tue in ML and programming, Poster Tue #118, Talk Wed in Deep reinforcement learning 1, Talk Wed in Language 2, Poster Wed #97, Poster Wed #82, Poster Wed #73
- Kolar, Mladen:** Talk Tue in Sparsity 2, Poster Tue #119
- Kolesnikov, Alexander:** Talk Mon in Deep generative models 1, Poster Mon #2
- Kolter, Zico:** Talk Mon in Continuous optimization 4, Talk Tue in Deep learning 6, Poster Tue #7, Poster Wed #5, Poster Wed #7
- Korolova, Aleksandra:** Workshop Fri in C4.5
- Kotlowski, Wojciech:** Talk Tue in Learning theory 2, Poster Tue #104
- Kottur, Satwik:** Talk Wed in Monte Carlo methods 1, Poster Wed #39
- Koutnik, Jan:** Talk Mon in Recurrent neural networks 1, Poster Mon #140
- Koyejo, Oluwasanmi:** Talk Tue in Learning theory 2, Poster Tue #104
- Krause, Andreas:** Talk Tue in Clustering 1, Talk Tue in Learning theory 2, Poster Tue #37, Poster Tue #76, Talk Wed in Privacy and security 1, Talk Wed in Combinatorial optimization 1, Talk Wed in Combinatorial optimization 2, Poster Wed #90, Poster Wed #106, Poster Wed #42, Poster Wed #72
- Krishnamurthy, Akshay:** Talk Mon in Active learning, Poster Mon #142, Talk Tue in Reinforcement learning 5, Poster Tue #65
- Kroer, Christian:** Talk Tue in Game theory and multiagents, Poster Tue #108
- Krompass, Denis:** Talk Tue in Recurrent neural networks 4, Poster Tue #122
- Krueger, David:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Kshirsagar, Meghana:** Workshop Thu in C4.4
- Kucukelbir, Alp:** Talk Mon in Probabilistic inference 1, Poster Mon #5, Poster Mon #21
- Kumar, Abhimanu:** Talk Tue in Probabilistic inference 3, Poster Tue #141
- Kumar, Sanjiv:** Talk Tue in Information theory, Poster Tue #112, Talk Wed in Large scale learning, Poster Wed #119
- Kumar, Ashish:** Talk Tue in Metric learning, Poster Wed #16
- Kumar, Naveen:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Kumar, Ravi:** Talk Mon in Matrix factorization 3, Poster Tue #29
- Kumar, Ashish:** Talk Mon in Supervised learning 1, Poster Mon #24
- Kumror, Daniel:** Talk Wed in Causal Inference 1, Poster Wed #93
- Kundaje, Anshul:** Talk Mon in Deep learning 1: backprop, Poster Mon #25
- Kurakin, Alexey:** Talk Wed in Large scale learning, Poster Wed #137
- Kushman, Nate:** Talk Tue in ML and programming, Poster Tue #134
- Kusner, Matt J.:** Talk Wed in Language 3, Poster Wed #132
- Kuznetsov, Vitaly:** Talk Mon in Deep learning 3: metalearning, Poster Mon #121, Workshop Fri in C4.1
- Kveton, Branislav:** Talk Mon in Online learning 1, Poster Mon #22, Poster Mon #30
- Kwok, James:** Talk Mon in Deep learning theory 2, Poster Mon #74
- Laclau, Charlotte:** Talk Tue in Clustering 2, Poster Wed #6
- Lacoste-Julien, Simon:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Lakshmanan, Laks:** Talk Mon in Online learning 1, Poster Mon #22
- Lakshminarayanan, Balaji:** Workshop Thu in Parkside 1
- Lamb, Alex:** Workshop Fri in C4.10
- Lampert, Christoph:** Talk Mon in Deep generative models 1, Talk Mon in Transfer and multitask learning, Poster Mon #2, Poster Mon #139
- Lang, Harry:** Talk Tue in Clustering 2, Poster Wed #12
- Lange, Jan-Hendrik:** Talk Wed in Combinatorial optimization 2, Poster Wed #115
- Langford, John:** Tutorial Sun in Tutorials Session A, Talk Mon in Supervised learning 2, Talk Mon in Active learning, Poster Mon #75, Poster Mon #142, Talk Tue in Reinforcement learning 5, Poster Tue #65
- Langseth, Helge:** Talk Mon in Probabilistic inference 1, Poster Mon #37
- Larsen, Rasmus:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Lattanzi, Silvio:** Talk Mon in Matrix factorization 3, Poster Tue #29
- Lattanzi, Silvio:** Talk Tue in Clustering 1, Poster Tue #45
- Laud, Peeter:** Workshop Fri in C4.5
- Law, Marc:** Talk Tue in Metric learning, Poster Wed #14
- Lawrence, Neil:** Talk Wed in Bayesian Optimization, Poster Wed #22
- Lazaric, Alessandro:** Talk Mon in Online learning 2, Talk Mon in Online learning 4, Poster Mon #81, Poster Tue #14
- Lazaro-Gredilla, Miguel:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Le, Hoang:** Talk Tue in Game theory and multiagents, Poster Tue #132
- Le, Quoc:** Talk Wed in Deep learning 8: hardware, Talk Wed in Deep reinforcement learning 2, Talk Wed in Large scale learning, Poster Wed #137, Poster Wed #134, Poster Wed #17
- Le, Hoai Minh:** Talk Mon in Continuous optimization 2, Poster Mon #60
- Le Thi, Hoai An:** Talk Mon in Continuous optimization 2, Poster Mon #60
- Leary, Lennox:** Talk Mon in Deep learning theory 1, Poster Tue #5
- LeCun, Yann:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Lee, Honglak:** Talk Tue in Recurrent neural networks 4, Poster Tue #106, Talk Wed in Deep reinforcement learning 1, Poster Wed #73
- Lee, Jungkwon:** Talk Mon in Deep generative models 2, Poster Mon #68
- Lee, Yin Tat:** Talk Wed in Distributed optimization, Poster Wed #37
- Lee, Kimin:** Talk Tue in Ensemble methods, Poster Tue #100
- Lee, Wee Sun:** Talk Tue in Probabilistic inference 1, Poster Tue #66
- Lee, Jun:** Talk Mon in Recurrent neural networks 1, Poster Mon #136
- Lee, Juho:** Talk Tue in Networks and relational learning, Poster Tue #43
- Lei, Qi:** Talk Mon in Continuous optimization 3, Poster Mon #105, Talk Tue in Information theory, Poster Tue #136
- Lei, Tao:** Talk Mon in Deep learning 2: invariances, Poster Mon #67
- Leme, Renato:** Talk Mon in Continuous optimization 1, Poster Mon #3
- Leon Suematsu, Yutaka:** Talk Wed in Large scale learning, Poster Wed #137
- Lerchner, Alexander:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Lessard, Laurent:** Talk Tue in Continuous optimization 6, Poster Tue #78
- Levine, Sergey:** Tutorial Sun in Tutorials Session B, Talk Mon in Reinforcement learning 2, Talk Mon in Deep learning 3: metalearning, Talk Mon in Reinforcement learning 4, Poster Mon #52, Poster Mon #112, Poster Tue #27, Talk Wed in Continuous control, Poster Wed #29
- Levy, Dor:** Talk Tue in ML and programming, Poster Tue #110
- Lewis, John:** Talk Mon in Deep learning theory 1, Poster Tue #5
- Lewis, Greg:** Talk Wed in Causal Inference 2, Poster Wed #127
- Leyton-Brown, Kevin:** Talk Wed in Causal Inference 2, Poster Wed #127
- Li, Hanbo:** Talk Tue in Ensemble methods, Poster Tue #93
- Li, Li Erran:** Workshop Thu in C4.10
- Li, Jialian:** Talk Mon in Online learning 2, Poster Mon #54
- Li, Lihong:** Talk Mon in Online learning 3, Poster Mon #108, Talk Tue in Reinforcement learning 5, Poster Tue #41
- Li, Xian:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Li, Cheng:** Talk Wed in Gaussian processes, Poster Wed #142
- Li, Yuanzhi:** Talk Mon in Online learning 2, Talk Mon in Matrix factorization 2, Poster Mon #100, Poster Mon #63, Talk Tue in Spectral methods, Poster Tue #44, Poster Tue #68, Talk Wed in Combinatorial optimization 2, Poster Wed #124
- Li, Chris Junchi:** Talk Mon in Latent feature models, Poster Mon #31
- Li, Jerry:** Talk Wed in High dimensional estimation, Talk Wed in Large scale learning, Poster Wed #23, Poster Wed #128
- Li, Hang:** Talk Wed in Language 1, Poster Wed #36
- Li, Jia:** Talk Wed in Monte Carlo methods 1, Poster Wed #48
- Li, Ke:** Talk Tue in Metric learning, Poster Wed #13
- Li, Xiangang:** Talk Wed in Language 1, Poster Wed #27
- Li, Shuai:** Talk Mon in Online learning 3, Poster Mon #99
- Li, Yingzhen:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #61
- Li, Wei-Chen:** Talk Mon in Latent feature models, Poster Mon #23
- Li, Qianxiao:** Talk Tue in Continuous optimization 6, Poster Tue #71
- Li, Qunwei:** Talk Tue in Continuous optimization 7, Poster Tue #131
- Li, Zina:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #102
- Li, Zhe:** Talk Mon in Deep learning theory 2, Poster Mon #83
- Li, Hao:** Talk Wed in Distributed optimization, Poster Wed #28
- Li, Shuang:** Talk Mon in Reinforcement learning 2, Poster Mon #79
- Li, Ping:** Talk Tue in Sparsity 1, Poster Tue #75
- Li, Jian:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #16
- Li, Chengtao:** Talk Mon in Bayesian Nonparametrics, Poster Mon #133
- Liang, Tengyuan:** Talk Mon in Online learning 4, Poster Tue #19
- Liang, Percy:** Talk Mon in Supervised learning 2, Talk Mon in Deep learning theory 3, Poster Mon #84, Poster Mon #101, Talk Tue in ML and programming, Poster Tue #142, Talk Wed in Applications, Poster Wed #113
- Liang, Yingbin:** Talk Tue in Continuous optimization 7, Poster Tue #131
- Liang, Xiaodan:** Talk Wed in Language 2, Poster Wed #80
- Liang, Yingyu:** Talk Mon in Deep generative models 2, Talk Mon in Matrix factorization 2, Poster Mon #100, Poster Mon #41, Talk Wed in Privacy and security 2, Poster Wed #86
- Liao, Siyu:** Talk Mon in Deep learning theory 2, Poster Mon #83
- Lillicrap, Timothy:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #6
- Lin, Shou-De:** Talk Mon in Latent feature models, Poster Mon #23
- Lin, Xunyu:** Talk Tue in Recurrent neural networks 4, Poster Tue #106
- Lin, Qihang:** Talk Mon in Continuous optimization 1, Talk Mon in Continuous optimization 3, Poster Mon #35, Poster Mon #114
- Lin, Di:** Talk Tue in Probabilistic inference 3, Poster Tue #125
- Lindgren, Erik:** Talk Tue in Probabilistic inference 1, Poster Tue #34
- Ling, Wang:** Workshop Thu in C4.11
- Ling, Jeffrey:** Talk Wed in Language 1, Poster Wed #45
- Lipor, John:** Talk Mon in Active learning, Poster Mon #134
- Lipton, Zachary:** Talk Wed in Applications, Poster Wed #104
- Littman, Michael:** Talk Mon in Reinforcement learning 2, Talk Mon in Reinforcement learning 3, Poster Mon #70, Poster Mon #97
- Liu, Tongliang:** Talk Tue in Learning theory 2, Poster Tue #97
- Liu, Ji:** Talk Tue in Sparsity 1, Poster Tue #89, Talk Wed in Large scale learning, Poster Wed #128
- Liu, Wei:** Talk Mon in Continuous optimization 2, Poster Mon #42, Talk Tue in Sparsity 2, Poster Tue #111
- Liu, Weiyang:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #120
- Liu, Han:** Talk Wed in High dimensional estimation, Poster Wed #50
- Liu, Hanxiao:** Talk Tue in Networks and relational learning, Poster Tue #51
- Liu, Hairong:** Talk Wed in Language 1, Poster Wed #27
- Liu, Hongwei:** Talk Wed in Monte Carlo methods 2, Poster Wed #100
- Liu, Jie:** Talk Tue in Continuous optimization 5, Poster Tue #48
- Liu, Shih-Chii:** Talk Mon in Recurrent neural networks 1, Poster Mon #136
- Liu, Liping:** Talk Mon in Deep generative models 3, Poster Mon #113
- Liu, Bo:** Talk Tue in Sparsity 1, Poster Tue #82



- Liu, Qingshan:** Talk Tue in Sparsity 1, Poster Tue #82
- Liu, Tie-Yan:** Talk Mon in Supervised learning 2, Poster Mon #48, Talk Wed in Distributed optimization, Poster Wed #19
- Liu, Peter:** Talk Tue in Recurrent neural networks 3, Poster Tue #70
- Liu, Yan:** Tutorial Sun in Tutorials Session C, Workshop Fri in C4.1
- Livescu, Karen:** Workshop Fri in C4.3
- Livni, Roi:** Talk Mon in Supervised learning 2, Poster Mon #57
- Loftin, Robert:** Talk Mon in Reinforcement learning 3, Poster Mon #97
- Long, Mingsheng:** Talk Mon in Deep learning 3: metalearning, Poster Mon #85
- Lorenzo, Orecchia:** Talk Mon in Continuous optimization 2, Poster Mon #78
- Lou, Xinghua:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Louizos, Christos:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #70
- Loukas, Andreas:** Talk Tue in Spectral methods, Poster Tue #36
- Low, Bryan Kian Hsiang:** Talk Wed in Gaussian processes, Poster Wed #108
- Lozano, Aurelie:** Talk Tue in Sparsity 1, Poster Tue #96
- Lu, Chun-Ta:** Talk Mon in Supervised learning 1, Poster Mon #39
- Lu, Zhengdong:** Talk Wed in Language 1, Poster Wed #36
- Lu, Yu:** Talk Mon in Online learning 3, Poster Mon #108
- Lu, Chi-Jen:** Talk Mon in Matrix factorization 3, Poster Tue #15
- Lu, Liang:** Workshop Fri in C4.3
- Lucchi, Aurelien:** Talk Tue in Continuous optimization 7, Poster Tue #107
- Lucey, Patrick:** Talk Tue in Game theory and multiagents, Poster Tue #132
- Lucic, Mario:** Talk Tue in Clustering 1, Talk Tue in Learning theory 2, Poster Tue #37, Poster Tue #76
- Lugosi, Gábor:** Talk Tue in Learning theory 2, Poster Tue #97
- Luo, Dixin:** Talk Wed in Time series, Poster Wed #78
- Luo, Ping:** Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #38
- Luo, Luo:** Talk Mon in Continuous optimization 4, Poster Tue #12
- Luong, Thang:** Talk Tue in Recurrent neural networks 3, Poster Tue #70
- Lv, Kaifeng:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #16
- Lyu, Michael:** Talk Wed in Robust Estimation, Poster Wed #67
- LYU, Yueming:** Talk Wed in Kernel methods, Poster Wed #138
- Ma, Wan-Duo:** Talk Mon in Deep learning theory 1, Poster Tue #5
- Ma, Guixiang:** Talk Mon in Supervised learning 1, Poster Mon #39
- Ma, Zhiming:** Talk Wed in Distributed optimization, Poster Wed #19
- Ma, Tengyu:** Talk Mon in Deep generative models 2, Poster Mon #41
- Ma, Shiqian:** Talk Mon in Continuous optimization 2, Poster Mon #42
- Ma, Yi-An:** Talk Wed in Monte Carlo methods 2, Poster Wed #92
- Ma, Fan:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #102
- Ma, Shuming:** Talk Mon in Deep learning 1: backprop, Poster Mon #17
- MacGlashan, James:** Talk Mon in Reinforcement learning 3, Poster Mon #97
- Mackey, Lester:** Talk Tue in Probabilistic inference 2, Poster Tue #101, Talk Wed in Monte Carlo methods 1, Poster Wed #56
- Madsen, Anders:** Talk Mon in Probabilistic learning 1, Poster Mon #37
- Mahajan, Dhruv:** Talk Tue in Ensemble methods, Poster Tue #79
- Maharaj, Tegan:** Talk Tue in Deep learning 7: analysis, Poster Tue #105
- Maheswaranathan, Niru:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #11
- Mahoney, Michael:** Talk Tue in Learning theory 1, Talk Tue in Spectral methods, Poster Tue #60, Poster Wed #1
- Maillard, Odalric:** Talk Tue in Spectral methods, Poster Tue #52
- Mair, Sebastian:** Talk Mon in Matrix factorization 1, Poster Mon #82
- Malherbe, Cédric:** Talk Mon in Continuous optimization 1, Poster Mon #19
- Malik, Jitendra:** Talk Tue in Metric learning, Poster Wed #13
- Malioutov, Dmitry:** Workshop Thu in C4.8
- Malkomes, Luiz Gustavo:** Talk Mon in Active learning, Poster Mon #130
- Mandt, Stephan:** Talk Wed in Language 1, Poster Wed #18
- Mankowitz, Daniel:** Workshop Thu in C4.1
- Mannor, Shie:** Talk Mon in Reinforcement learning 3, Poster Mon #106, Talk Tue in Reinforcement learning 5, Poster Tue #57, Workshop Thu in C4.1
- Mannor, Shie:** Talk Mon in Online learning 1, Poster Mon #6
- Mao, Xueyu:** Talk Mon in Matrix factorization 1, Poster Mon #64
- Marks, Debora:** Talk Tue in Probabilistic inference 1, Poster Tue #58
- Martin, Andrew:** Talk Tue in Ensemble methods, Poster Tue #93
- Masegosa, Andres:** Talk Mon in Probabilistic learning 1, Poster Mon #37
- Massoulié, Laurent:** Talk Wed in Distributed optimization, Poster Wed #37
- Mastromatteo, Iacopo:** Talk Wed in Causal Inference 1, Poster Wed #75
- Matei, Basarab:** Talk Tue in Clustering 2, Poster Wed #6
- Matsen, Frederick:** Talk Wed in Monte Carlo methods 2, Poster Wed #74
- Matsumoto, Eiichi:** Talk Tue in Information theory, Poster Tue #144
- Matthey, Loic:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Maturana, Daniel:** Talk Wed in Continuous control, Poster Wed #20
- Matveev, Alexander:** Talk Wed in Deep learning 8: hardware, Poster Wed #26
- Maystre, Lucas:** Talk Tue in Ranking and preferences, Poster Tue #74, Poster Tue #88
- Mazumdar, Arya:** Talk Wed in High dimensional estimation, Poster Wed #32
- McAuley, Julian:** Talk Wed in Applications, Poster Wed #104
- McCallum, Andrew:** Talk Wed in Structured prediction, Poster Wed #130
- McGill, Mason:** Talk Wed in Deep reinforcement learning 2, Poster Wed #116
- McKenna, Ryan:** Talk Wed in Privacy and security 1, Poster Wed #51
- McMahan, H. Brendan:** Talk Tue in Information theory, Poster Tue #112
- McNamara, Daniel:** Talk Mon in Transfer and multitask learning, Poster Mon #146
- McWilliams, Brian:** Talk Mon in Deep learning theory 1, Poster Tue #2, Poster Tue #5
- Mei, Jiali:** Talk Mon in Matrix factorization 1, Poster Mon #73
- Meng, Deyu:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #102
- Meng, Qi:** Talk Wed in Distributed optimization, Poster Wed #19
- Menick, Jacob:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #127
- Mescheder, Lars:** Talk Mon in Deep generative models 1, Poster Mon #26
- Mesgarani, Nima:** Talk Tue in Deep learning 7: analysis, Poster Tue #121
- Meshi, Ofer:** Workshop Fri in C4.5
- Metaxas, Dimitris:** Talk Tue in Sparsity 1, Poster Tue #82
- mhammedi, zakaria:** Talk Tue in Recurrent neural networks 2, Poster Tue #31
- Miao, Yishu:** Talk Wed in Language 2, Poster Wed #62, Poster Wed #71, Workshop Thu in C4.11, Workshop Fri in C4.5
- Michiardi, Pietro:** Talk Wed in Gaussian processes, Poster Wed #126
- Miklau, Gerome:** Talk Wed in Privacy and security 1, Poster Wed #51
- Miller, Andrew:** Talk Tue in Probabilistic inference 3, Poster Tue #109
- Miller, John:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Milli, Smitha:** Workshop Fri in C4.7
- Mineiro, Paul:** Talk Mon in Supervised learning 2, Poster Mon #75
- Mirhoseini, Azalia:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Mirroknj, Vahab:** Talk Mon in Continuous optimization 1, Poster Mon #3
- Mirzasoleiman, Baharan:** Talk Wed in Combinatorial optimization 2, Poster Wed #106
- Mishra, Nikhil:** Talk Mon in Reinforcement learning 2, Poster Mon #61
- Mitliagkas, Ioannis:** Talk Wed in Monte Carlo methods 1, Poster Wed #56
- Mitrovic, Marko:** Talk Wed in Privacy and security 1, Poster Wed #42, Workshop Fri in C4.5
- Mitrovic, Slobodan:** Talk Wed in Combinatorial optimization 1, Poster Wed #81
- Miyato, Takeru:** Talk Tue in Information theory, Poster Tue #144
- Modhe, Nirbhay:** Talk Wed in Structured prediction, Poster Wed #144
- Mohajer, Soheil:** Talk Tue in Ranking and preferences, Poster Tue #102
- Mohamed, Shakir:** Workshop Thu in Parkside 1
- Mohammad, Abdelrahman:** Talk Tue in ML and programming, Talk Tue in Recurrent neural networks 4, Poster Tue #130, Poster Tue #118
- Mohri, Mehryar:** Talk Mon in Deep learning 3: metalearning, Poster Mon #121
- Moitra, Ankur:** Tutorial Sun in Tutorials Session C, Talk Mon in Probabilistic learning 2, Poster Mon #62, Talk Wed in High dimensional estimation, Poster Wed #23
- Molchanov, Dmitry:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #79
- Mollaysa, Amina:** Talk Tue in Deep learning 6, Poster Wed #11
- Moore, Sherry:** Talk Wed in Large scale learning, Poster Wed #137
- Mordatch, Igor:** Talk Mon in Reinforcement learning 2, Poster Mon #61
- Moreno, Alexander:** Talk Wed in Healthcare, Poster Wed #52
- Morgenstern, Jamie:** Talk Mon in Reinforcement learning 1, Poster Mon #20
- Morrison, Clayton:** Talk Mon in Bayesian Nonparametrics, Poster Mon #129
- Moseley, Benjamin:** Talk Mon in Active learning, Poster Mon #130
- Mou, Wenlong:** Talk Wed in Privacy and security 2, Poster Wed #86, Poster Wed #102
- Mou, Lili:** Talk Wed in Language 1, Poster Wed #36
- Moura, Jose:** Talk Wed in Monte Carlo methods 1, Poster Wed #39
- Mroueh, Youssef:** Talk Mon in Deep generative models 2, Poster Mon #50
- Mueller, Jonas:** Talk Tue in Recurrent neural networks 4, Poster Tue #114
- Mueller, Klaus-robort:** Talk Wed in Privacy and security 1, Poster Wed #59
- Mukherjee, Sayan:** Talk Mon in Matrix factorization 1, Poster Mon #55
- Mukkamala, Mahesh Chandra:** Talk Mon in Online learning 4, Poster Tue #28
- Munkhdalai, Tsendsuren:** Talk Mon in Deep learning 3: metalearning, Poster Mon #94
- Munos, Remi:** Talk Mon in Reinforcement learning 1, Talk Mon in Semisupervised and curriculum learning, Talk Mon in Reinforcement learning 4, Poster Mon #127, Poster Mon #12, Poster Tue #13, Talk Wed in Deep reinforcement learning 1, Poster Wed #64
- Murphy, Susan:** Talk Wed in Healthcare, Poster Wed #52
- Musco, Cameron:** Talk Mon in Supervised learning 2, Poster Mon #66
- Musco, Christopher:** Talk Mon in Supervised learning 2, Poster Mon #66
- Muzu, Jean-François:** Talk Wed in Causal Inference 1, Poster Wed #75
- Mylymaki, Petri:** Workshop Fri in C4.11
- Mély, David:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Nagamine, Tasha:** Talk Tue in Deep learning 7: analysis, Poster Tue #121
- Nakagawa, Kazuya:** Talk Tue in Sparsity 2, Poster Tue #135
- Nakahara, Hiroyuki:** Talk Mon in Continuous optimization 4, Poster Tue #26
- Nakajima, Shinichi:** Talk Wed in Privacy and security 1, Poster Wed #59
- Namkoong, Hongseok:** Talk Tue in Continuous optimization 6, Poster Tue #99
- Naradowsky, Jason:** Talk Tue in ML and programming, Poster Tue #126
- Nardelli, Nantas:** Talk Wed in Deep reinforcement learning 1, Poster Wed #82
- Natarajan, Nagarajan:** Talk Mon in Active learning, Poster Mon #138, Talk Tue in Learning theory 2, Poster Tue #104
- Negahban, Sahand:** Talk Wed in Combinatorial optimization 1, Poster Wed #98
- Neil, Daniel:** Talk Mon in Recurrent neural networks 1, Poster Mon #136
- Neiswanger, William:** Talk Mon in Probabilistic learning 1, Poster Mon #13
- Netrapalli, Praneeth:** Talk Tue in Continuous optimization 7, Poster Tue #139
- Neu, Gergely:** Talk Tue in Learning theory 2, Poster Tue #97



- Neumann, Gerhard:** Talk Wed in Continuous control, Poster Wed #55
- Ng, Andrew:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Nguyen, Vu:** Talk Wed in Gaussian processes, Poster Wed #142
- Nguyen, Lam:** Talk Tue in Continuous optimization 5, Poster Tue #48
- Nguyen, Quynh:** Talk Mon in Deep learning theory 1, Poster Tue #1
- Nguyen, Long:** Talk Tue in Clustering 1, Poster Tue #69
- NI, XIUYAN:** Talk Mon in Probabilistic learning 3, Poster Mon #116
- Niekum, Scott:** Talk Tue in Reinforcement learning 5, Poster Tue #33
- Nielsen, Frank:** Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #30
- Nielsen, Thomas:** Talk Mon in Probabilistic learning 1, Poster Mon #37
- Nielsen, Frank:** Talk Tue in Recurrent neural networks 3, Poster Tue #91
- Niu, Gang:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #111
- Nock, Richard:** Workshop Fri in C4.8
- Norouzi, Mohammad:** Talk Tue in Recurrent neural networks 3, Poster Tue #98, Talk Wed in Deep learning 8: hardware, Talk Wed in Structured prediction, Poster Wed #112, Poster Wed #17
- Nowak, Robert:** Talk Wed in Robust Estimation, Poster Wed #85
- Nowozin, Sebastian:** Talk Mon in Deep generative models 1, Poster Mon #26
- Oates, Chris:** Talk Tue in Probabilistic inference 2, Poster Tue #94
- Ochiai, Tsubasa:** Talk Wed in Language 1, Poster Wed #54
- Odena, Augustus:** Talk Mon in Deep generative models 2, Poster Mon #59
- Oglic, Dino:** Talk Wed in Kernel methods, Poster Wed #143
- Oh, Junhyuk:** Talk Wed in Deep reinforcement learning 1, Poster Wed #73
- Olah, Christopher:** Talk Mon in Deep generative models 2, Poster Mon #59
- Oliehoek, Frans:** Talk Mon in Reinforcement learning 3, Poster Mon #115
- Oliva, Junier:** Talk Tue in Recurrent neural networks 2, Poster Tue #55
- Omidshafiei, Shayegan:** Talk Tue in Game theory and multiagents, Poster Tue #140
- Ong, Cheng Soon:** Workshop Fri in C4.8
- ONG, yEW:** Talk Mon in Transfer and multitask learning, Poster Mon #135
- Ongie, Greg:** Talk Wed in Robust Estimation, Poster Wed #85
- Orabona, Francesco:** Talk Mon in Online learning 4, Poster Tue #9
- Orlitsky, Alon:** Talk Tue in Ranking and preferences, Talk Tue in Information theory, Poster Tue #128, Poster Tue #95
- Osband, Ian:** Talk Mon in Reinforcement learning 1, Poster Mon #36
- Osband, Ian:** Talk Mon in Reinforcement learning 1, Poster Mon #12
- Osindero, Simon:** Talk Mon in Deep learning 1: backprop, Poster Mon #1, Poster Mon #9, Talk Wed in Deep reinforcement learning 2, Poster Wed #107
- Osoyama, Takayuki:** Talk Wed in Time series, Poster Wed #69
- Ostrovski, Georg:** Talk Wed in Deep reinforcement learning 1, Poster Wed #64
- Oudot, Steve:** Talk Wed in Kernel methods, Poster Wed #120
- Ozer, Sedat:** Talk Mon in Matrix factorization 1, Poster Mon #46
- Pachet, François:** Talk Tue in Recurrent neural networks 3, Poster Tue #91
- Pad, Pedram:** Talk Tue in Sparsity 2, Poster Tue #143
- Paige, Brooks:** Talk Wed in Language 3, Poster Wed #132
- Pakman, Ari:** Talk Wed in Monte Carlo methods 1, Poster Wed #30
- Pal, Christopher:** Talk Tue in Recurrent neural networks 2, Poster Tue #39
- Pal, Christopher:** Talk Wed in Deep learning 9: probabilistic, Poster Wed #88
- Pal, Arka:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Pal, David:** Talk Mon in Online learning 4, Poster Tue #19
- Palla, Konstantina:** Talk Mon in Bayesian Nonparametrics, Poster Mon #144
- Pan, Yangchen:** Talk Wed in Kernel methods, Poster Wed #129
- Pan, Yunpeng:** Talk Wed in Continuous control, Poster Wed #38
- Panahi, Ashkan:** Talk Tue in Clustering 2, Poster Wed #10
- Panigrahy, Rina:** Talk Mon in Matrix factorization 3, Poster Tue #29
- Paninski, Liam:** Talk Wed in Monte Carlo methods 1, Poster Wed #30
- Paranjape, Bhargavi:** Talk Tue in Metric learning, Poster Wed #16
- Park, Yookoon:** Talk Mon in Deep learning 3: metalearning, Poster Mon #103
- Park, KyoungSoo:** Talk Tue in Ensemble methods, Poster Tue #100
- Park, Kyoungsoo:** Talk Tue in Probabilistic inference 2, Poster Tue #73
- Pascanu, Razvan:** Talk Mon in Deep learning theory 1, Poster Tue #3
- Pathak, Deepak:** Talk Mon in Reinforcement learning 3, Poster Mon #88
- Pazis, Jason:** Talk Tue in Game theory and multiagents, Poster Tue #140
- Pe'er, Dana:** Workshop Thu in C4.4
- Pearl, Judea:** Talk Wed in Causal Inference 2, Poster Wed #136
- Peng, Bei:** Talk Mon in Reinforcement learning 3, Poster Mon #97
- Peng, Hao:** Talk Wed in Gaussian processes, Poster Wed #135
- Pennington, Jeffrey:** Talk Mon in Deep learning theory 1, Poster Tue #4
- Pentina, Anastasia:** Talk Mon in Transfer and multitask learning, Poster Mon #139
- Perlin, Ken:** Talk Wed in Applications, Poster Wed #140
- Perona, Pietro:** Talk Tue in Ensemble methods, Poster Tue #72, Talk Wed in Deep reinforcement learning 2, Poster Wed #116
- Peters, Jan:** Talk Wed in Continuous control, Poster Wed #55
- Peurifoy, John:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Pham, Hieu:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Phan, Duy Nhat:** Talk Mon in Continuous optimization 2, Poster Mon #60
- Phoenix, Scott:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Phung, Dinh:** Talk Tue in Clustering 1, Poster Tue #69
- Pichapati, Venkatadheeraj:** Talk Tue in Ranking and preferences, Poster Tue #95
- Pineau, Joelle:** Workshop Fri in C4.10
- Pinto, Lrel:** Talk Mon in Reinforcement learning 1, Poster Mon #4
- Pirotta, Matteo:** Talk Mon in Reinforcement learning 1, Poster Mon #28
- Pleiss, Geoff:** Talk Tue in Deep learning 7: analysis, Poster Tue #137
- Pokutta, Sebastian:** Talk Mon in Online learning 4, Talk Tue in Continuous optimization 5, Poster Tue #32, Poster Tue #40, Poster Tue #24
- Pontil, Massimiliano:** Talk Tue in Continuous optimization 6, Poster Tue #92
- Poole, Ben:** Talk Mon in Deep learning theory 3, Poster Mon #110, Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #46
- Prabhakaran, Sandhya:** Workshop Thu in C4.4
- Precup, Doina:** Workshop Fri in Parkside 1
- Price, Eric:** Talk Tue in Sparsity 1, Poster Tue #103
- Pritzel, Alexander:** Talk Mon in Reinforcement learning 3, Poster Mon #123, Talk Wed in Deep reinforcement learning 2, Poster Wed #125
- Pronobis, Andrzej:** Workshop Thu in C4.5
- Puigdomenech Badia, Adrià:** Talk Wed in Deep reinforcement learning 2, Poster Wed #125
- Pyzer-Knapp, Edward:** Talk Wed in Bayesian Optimization, Poster Wed #57
- Póczos, Barnabás:** Talk Mon in Deep learning 2: invariances, Poster Mon #40, Talk Tue in Recurrent neural networks 2, Talk Tue in Information theory, Poster Tue #120, Poster Tue #55, Talk Wed in Bayesian Optimization, Poster Wed #49
- Qi, Yuan:** Talk Wed in Gaussian processes, Poster Wed #135
- Qi, Guo-Jun:** Talk Mon in Recurrent neural networks 1, Poster Mon #132
- Qin, Tao:** Talk Mon in Supervised learning 2, Poster Mon #48
- Quadrianto, Novi:** Talk Mon in Probabilistic learning 3, Poster Mon #116
- Rabinowitz, Neil:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- Raffel, Colin:** Talk Tue in Recurrent neural networks 3, Poster Tue #70
- Raghu, Maithra:** Talk Mon in Deep learning theory 3, Poster Mon #110
- Raghunathan, Aditi:** Talk Tue in Learning theory 1, Poster Wed #2
- Ragin, Ann:** Talk Mon in Supervised learning 1, Poster Mon #39
- Ragni, Anton:** Workshop Fri in C4.3
- Rahman, Ashfaqur:** Talk Tue in Recurrent neural networks 2, Poster Tue #31
- Rahmani, Mostafa:** Talk Mon in Matrix factorization 2, Poster Mon #118, Talk Tue in Sparsity 2, Poster Tue #127
- Rai, Piyush:** Talk Tue in Networks and relational learning, Poster Tue #67, Talk Wed in Structured prediction, Poster Wed #144
- Raiman, Jonathan:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Raj, Anant:** Talk Tue in Continuous optimization 5, Poster Tue #56
- Ramchandran, Kannan:** Talk Mon in Online learning 1, Poster Mon #38
- Ramos-Lopez, Dario:** Talk Mon in Probabilistic learning 1, Poster Mon #37
- Rana, Santu:** Talk Wed in Gaussian processes, Poster Wed #142
- Ranganath, Rajesh:** Workshop Thu in Parkside 1
- Rao, Satish:** Talk Tue in Spectral methods, Poster Tue #60
- Ratner, Alexander:** Talk Mon in Probabilistic learning 3, Poster Mon #89
- Ravanbakhsh, Siamak:** Talk Mon in Deep learning 2: invariances, Poster Mon #40
- Ravikumar, Pradeep:** Talk Mon in Latent feature models, Talk Mon in Probabilistic learning 2, Talk Mon in Continuous optimization 3, Poster Mon #44, Poster Mon #105, Poster Mon #23
- Ravindran, Balaraman:** Workshop Thu in C4.1
- Ray Chaudhuri, Shraman:** Talk Wed in Deep learning 8: hardware, Poster Wed #26
- Ray Chowdhury, Sayak:** Talk Mon in Online learning 2, Poster Mon #72
- Re, Christopher:** Talk Mon in Probabilistic learning 3, Poster Mon #89
- Real, Esteban:** Talk Wed in Large scale learning, Poster Wed #137
- Recht, Benjamin:** Talk Mon in Continuous optimization 2, Poster Mon #51
- Redko, Ievgen:** Talk Tue in Clustering 2, Poster Wed #6
- Reed, Scott:** Talk Mon in Deep generative models 1, Poster Mon #10
- Rehg, Jim:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #120, Talk Wed in Healthcare, Poster Wed #52
- Reichert, David:** Talk Wed in Deep reinforcement learning 1, Poster Wed #91
- REN, Xuancheng:** Talk Mon in Deep learning 1: backprop, Poster Mon #17
- Requeima, James:** Talk Wed in Bayesian Optimization, Poster Wed #57
- Resnick, Cinjon:** Talk Tue in Recurrent neural networks 3, Poster Tue #98
- Restelli, Marcello:** Talk Mon in Reinforcement learning 1, Poster Mon #28
- Rich, Zemel:** Talk Tue in Metric learning, Poster Wed #14
- Riedel, Sebastian:** Talk Tue in ML and programming, Poster Tue #126
- Rigollet, Philippe:** Talk Mon in Probabilistic learning 2, Poster Mon #62
- Riley, Patrick:** Talk Wed in Applications, Poster Wed #131
- Rippel, Oren:** Talk Wed in Applications, Poster Wed #122
- Riquelme Ruiz, Carlos:** Talk Mon in Online learning 4, Poster Tue #14
- Ritter, Julian Hippolyt:** Talk Mon in Continuous optimization 4, Poster Tue #22
- Ritter, Samuel:** Talk Tue in Deep learning 7: analysis, Poster Tue #113
- Roberts, David:** Talk Mon in Reinforcement learning 3, Poster Mon #97
- Roberts, Adam:** Talk Tue in Recurrent neural networks 3, Poster Tue #98
- Rocktäschel, Tim:** Talk Tue in ML and programming, Poster Tue #126
- Rosman, Benjamin:** Talk Mon in Reinforcement learning 4, Poster Tue #18
- Roth, Aaron:** Talk Mon in Reinforcement learning 1, Poster Mon #20, Talk Tue in Learning theory 1, Poster Wed #3
- Rowland, Mark:** Talk Wed in Monte Carlo methods 2, Poster Wed #65
- Rubinstein, Benjamin:** Talk Wed in Privacy and security 2, Poster Wed #77
- Rudin, Cynthia:** Talk Mon in Probabilistic learning 2, Poster Mon #53



- Ruggieri, Salvatore:** Talk Mon in Supervised learning 1, Poster Mon #8
- Rukat, Tammo:** Talk Mon in Probabilistic learning 2, Poster Mon #80
- Rus, Daniela:** Talk Mon in Matrix factorization 1, Poster Mon #46
- Rush, Alexander:** Talk Wed in Language 1, Poster Wed #45
- Rusu, Andrei:** Talk Mon in Reinforcement learning 3, Poster Mon #123
- Safran, Itay:** Talk Mon in Deep learning theory 3, Poster Mon #119
- Sagarna, Ramon:** Talk Mon in Transfer and multitask learning, Poster Mon #135
- Sainath, Tara:** Workshop Fri in C4.3
- Saito, Kuniaki:** Talk Mon in Transfer and multitask learning, Poster Mon #131
- Sakai, Tomoya:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #111
- Sakr, Charbel:** Talk Mon in Deep learning theory 2, Poster Mon #65
- Sakuma, Jun:** Talk Wed in Privacy and security 2, Poster Wed #95
- Salakhutdinov, Ruslan:** Talk Wed in Language 2, Talk Wed in Language 3, Poster Wed #80, Poster Wed #123
- Salehi, Farnood:** Talk Tue in Sparsity 2, Poster Tue #143
- Saligrama, Venkatesh:** Talk Mon in Continuous optimization 2, Poster Mon #78, Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #54, Workshop Thu in C4.7
- Salmeron, Antonio:** Talk Mon in Probabilistic learning 1, Poster Mon #37
- Salzmann, Mathieu:** Talk Tue in Metric learning, Poster Wed #15
- Sandholm, Tuomas:** Talk Tue in Game theory and multiagents, Poster Tue #108, Poster Tue #116
- Santoro, Adam:** Talk Tue in Deep learning 7: analysis, Poster Tue #113
- Santurkar, Shibani:** Talk Wed in Deep learning 8: hardware, Poster Tue #26
- Sarkar, Purnamrita:** Talk Mon in Matrix factorization 1, Poster Mon #64
- Satheesh, Sanjeev:** Talk Wed in Language 1, Poster Wed #27
- Sato, Issei:** Talk Mon in Deep learning 1: backprop, Poster Mon #33
- Savarese, Silvio:** Workshop Thu in C4.10
- Saxe, Andrew:** Talk Mon in Reinforcement learning 4, Poster Tue #18
- Saxena, Saurabh:** Talk Wed in Large scale learning, Poster Wed #137
- Scaman, Kevin:** Talk Wed in Distributed optimization, Poster Wed #37
- Scarlett, Jonathan:** Talk Wed in Combinatorial optimization 1, Poster Wed #81
- Schaal, Stefan:** Talk Wed in Continuous control, Poster Wed #29
- Schapiro, Robert:** Talk Tue in Reinforcement learning 5, Poster Tue #65
- Schaul, Tom:** Talk Wed in Deep reinforcement learning 1, Talk Wed in Deep reinforcement learning 2, Poster Wed #107, Poster Wed #91
- Schaul, Tom:** Workshop Thu in C4.6
- Scheinberg, Katya:** Talk Tue in Continuous optimization 5, Poster Tue #48
- Scherer, Sebastian:** Talk Wed in Continuous control, Poster Wed #20
- Schlachter, Kristofer:** Talk Wed in Applications, Poster Wed #140
- Schlegel, Matthew:** Talk Wed in Kernel methods, Poster Wed #129
- Schmidhuber, Jürgen:** Talk Mon in Recurrent neural networks 1, Poster Mon #140
- Schmidt, Erik:** Workshop Fri in C4.9
- Schmidt, Mark:** Talk Mon in Online learning 1, Poster Mon #22
- Schneider, Jeff:** Talk Tue in Recurrent neural networks 2, Poster Tue #55, Talk Wed in Bayesian Optimization, Poster Wed #49
- Schneider, Jeff:** Talk Mon in Deep learning 2: invariances, Poster Mon #40
- Schneider, Oskar:** Talk Mon in Online learning 4, Poster Tue #24
- Schoenholz, Samuel:** Talk Wed in Applications, Poster Wed #131
- Schwing, Alex:** Workshop Fri in C4.5
- Schölkopf, Prof. Bernhard:** Invited Talk Wed in Darling Harbour Theater
- Scott, Chris:** Talk Wed in Healthcare, Poster Wed #52
- Seeger, Matthias:** Talk Wed in Bayesian Optimization, Poster Wed #40
- Sekiyama, Taro:** Talk Wed in Time series, Poster Wed #69
- Selle, Andrew:** Talk Wed in Large scale learning, Poster Wed #137
- Selsam, Daniel:** Talk Tue in ML and programming, Poster Tue #142
- Seltzer, Margo:** Talk Mon in Probabilistic learning 2, Poster Mon #53
- Sen, Rajat:** Talk Wed in Causal Inference 2, Poster Wed #118
- Sengupta, Shubho:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Sercu, Tom:** Talk Mon in Deep generative models 2, Poster Mon #50
- Shakkottai, Sanjay:** Talk Wed in Causal Inference 2, Poster Wed #118
- Shalev-Shwartz, Shai:** Talk Mon in Deep learning theory 2, Poster Mon #56
- Shalit, Uri:** Talk Wed in Causal Inference 1, Poster Wed #101
- Shamir, Ohad:** Talk Mon in Latent feature models, Talk Mon in Continuous optimization 1, Talk Mon in Online learning 1, Talk Mon in Deep learning theory 2, Talk Mon in Deep learning theory 3, Poster Mon #11, Poster Mon #14, Poster Mon #7, Poster Mon #119, Poster Mon #56
- Shammah, Shaked:** Talk Mon in Deep learning theory 2, Poster Mon #56
- Shanbhag, Naresh:** Talk Mon in Deep learning theory 2, Poster Mon #65
- Shanmugam, Karthikeyan:** Talk Wed in Causal Inference 2, Poster Wed #118
- Sharan, Vatsal:** Talk Mon in Matrix factorization 3, Poster Tue #10
- Shavit, Nir:** Talk Wed in Deep learning 8: hardware, Poster Wed #26
- Sheffet, Or:** Talk Wed in Privacy and security 1, Poster Wed #24
- Sheldon, Daniel:** Talk Tue in Probabilistic inference 1, Poster Tue #42, Talk Wed in Privacy and security 1, Poster Wed #51
- Shen, Yichen:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Shen, Dinghan:** Talk Wed in Language 2, Poster Wed #89
- Shen, Jie:** Talk Tue in Sparsity 1, Poster Tue #75
- Shen, Linlin:** Talk Mon in Supervised learning 1, Poster Mon #39
- Shen, Li:** Talk Mon in Continuous optimization 2, Poster Mon #42
- Shi, Tim:** Talk Wed in Applications, Poster Wed #113
- Shimizu, Kana:** Workshop Fri in C4.4
- Shimkin, Nahum:** Talk Wed in Deep reinforcement learning 1, Poster Wed #99
- Shin, Jinwoo:** Talk Tue in Probabilistic inference 2, Talk Tue in Ensemble methods, Poster Tue #73, Poster Tue #100
- Shlens, Jon:** Talk Mon in Deep generative models 2, Poster Mon #59
- Shoeybi, Mohammad:** Talk Tue in Recurrent neural networks 3, Poster Tue #84
- Shofner, Alyssa:** Talk Mon in Active learning, Poster Mon #130
- Shrikumar, Avanti:** Talk Mon in Deep learning 1: backprop, Poster Mon #25
- Shrivastava, Anshumali:** Talk Wed in Large scale learning, Poster Wed #110
- Shu, Rui:** Talk Mon in Deep generative models 3, Poster Mon #95
- Shyam, Pranav:** Talk Mon in Recurrent neural networks 1, Poster Mon #128
- Si, Si:** Talk Tue in Ensemble methods, Poster Tue #79
- Sidford, Aaron:** Talk Tue in Continuous optimization 7, Poster Tue #123
- Sidiropoulos, Nicholas:** Talk Tue in Clustering 1, Poster Tue #53
- Sidor, Szymon:** Talk Mon in Transfer and multitask learning, Poster Mon #143
- Silver, David:** Talk Mon in Deep learning 1: backprop, Poster Mon #1, Talk Wed in Deep reinforcement learning 1, Talk Wed in Deep reinforcement learning 2, Poster Wed #107, Poster Wed #91
- Silverman, Edwin:** Talk Tue in Clustering 2, Poster Wed #8
- Simhadri, Harsha Vardhan:** Talk Tue in Metric learning, Poster Wed #16
- Simonyan, Karen:** Talk Mon in Deep generative models 1, Poster Mon #18, Talk Tue in Recurrent neural networks 3, Poster Tue #98
- Simsekli, Umut:** Talk Wed in Monte Carlo methods 1, Poster Wed #21
- Singer, Yaron:** Talk Wed in Combinatorial optimization 1, Poster Wed #63
- Singh, Rishabh:** Talk Tue in ML and programming, Poster Tue #118
- Singh, Karan:** Talk Mon in Online learning 2, Poster Mon #45, Talk Wed in Privacy and security 2, Poster Wed #68
- Singh, Vikas:** Talk Wed in Healthcare, Poster Wed #25
- Singh, Satinder:** Talk Wed in Deep reinforcement learning 1, Poster Wed #73
- Singh, Aarti:** Talk Tue in Probabilistic inference 3, Poster Tue #133, Talk Wed in Combinatorial optimization 2, Poster Wed #124
- Singh, Shashank:** Talk Tue in Information theory, Poster Tue #120
- Sinha, Aman:** Talk Tue in Continuous optimization 6, Poster Tue #99
- Sivakumar, Vidyashankar:** Talk Wed in High dimensional estimation, Poster Wed #41
- Skirlo, Scott:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Smith, Linda:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #120
- Smola, Alex:** Tutorial Sun in Tutorials Session A, Talk Tue in Recurrent neural networks 4, Poster Tue #138, Talk Wed in Monte Carlo methods 1, Poster Wed #39
- Sohl-Dickstein, Jascha:** Talk Mon in Deep learning theory 3, Talk Mon in Deep learning 4: learning to learn, Poster Mon #110, Talk Tue in Recurrent neural networks 2, Poster Tue #63, Poster Tue #11
- Sohler, Christian:** Talk Tue in Clustering 2, Poster Wed #12
- Sohn, Sungryull:** Talk Tue in Recurrent neural networks 4, Poster Tue #106
- Solja, Marin:** Talk Tue in Recurrent neural networks 2, Poster Tue #47
- Song, Zhao:** Talk Mon in Deep learning theory 2, Poster Mon #47
- Song, Le:** Talk Mon in Reinforcement learning 2, Talk Mon in Semisupervised and curriculum learning, Poster Mon #79, Poster Mon #120, Talk Tue in Networks and relational learning, Poster Tue #59, Talk Wed in Time series, Talk Wed in Large scale learning, Poster Wed #103, Poster Wed #119
- Song, Jiaming:** Talk Mon in Deep generative models 3, Poster Mon #86
- Sontag, David:** Talk Mon in Supervised learning 1, Poster Mon #16, Talk Wed in Causal Inference 1, Poster Wed #101
- Sordoni, Alessandro:** Talk Mon in Deep learning 4: learning to learn, Poster Tue #21
- Sorokin, Dmitry:** Talk Wed in Continuous control, Poster Wed #55
- Sprechmann, Pablo:** Talk Wed in Applications, Poster Wed #140
- Srebro, Nati:** Talk Mon in Latent feature models, Poster Mon #7, Talk Tue in Sparsity 2, Poster Tue #119
- Sriram, Srinivasan:** Talk Wed in Deep reinforcement learning 2, Poster Wed #125
- Srivastava, Rupesh:** Talk Mon in Recurrent neural networks 1, Poster Mon #140
- Stab, Matthew:** Talk Wed in Combinatorial optimization 2, Poster Wed #133
- Stan, Serban:** Talk Wed in Combinatorial optimization 1, Poster Wed #90
- Stefankovic, Daniel:** Talk Tue in Sparsity 1, Poster Tue #89
- steiner, benoit:** Talk Wed in Deep learning 8: hardware, Poster Wed #17
- Steinhardt, Jacob:** Workshop Fri in C4.7
- Stewart, Alistair:** Talk Wed in High dimensional estimation, Poster Wed #23
- Stich, Sebastian:** Talk Tue in Continuous optimization 5, Poster Tue #56
- Stone, Peter:** Talk Tue in Reinforcement learning 5, Poster Tue #33
- Strasser, Pablo:** Talk Tue in Deep learning 6, Poster Wed #11
- Studer, Christoph:** Talk Mon in Continuous optimization 3, Poster Mon #122
- Sudderth, Erik:** Talk Mon in Bayesian Nonparametrics, Poster Mon #137
- SUGGALA, ARUN SAI:** Talk Mon in Probabilistic learning 2, Poster Mon #44
- SUGGALA, ARUN:** Talk Tue in Metric learning, Poster Tue #16
- Suggala, Arun:** Talk Mon in Latent feature models, Poster Mon #23
- Sugiyama, Mahito:** Talk Mon in Continuous optimization 4, Poster Tue #26
- Sugiyama, Masashi:** Talk Mon in Semisupervised and curriculum learning, Poster Mon #111, Talk Tue in Information theory, Poster Tue #144
- Suh, Changho:** Talk Tue in Ranking and preferences, Poster Tue #102
- Sujono, Debora:** Talk Tue in Probabilistic inference 1, Poster Tue #42
- Sukhatme, Gaurav:** Talk Wed in Continuous control, Poster Wed #29
- Sukthankar, Rahul:** Talk Mon in Reinforcement learning 1, Poster Mon #4
- Sun, Yu:** Talk Tue in Deep learning 7: analysis, Poster Tue #137

Sun, Tao: Talk Wed in Privacy and security 1, Poster Wed #51

Sun, Wen: Talk Mon in Online learning 3, Poster Mon #117, Talk Wed in Structured prediction, Poster Wed #121

Sun, Ke: Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #30

Sun, Jimeng: Tutorial Sun in Tutorials Session C

SUN, Xu: Talk Mon in Deep learning 1: backprop, Poster Mon #17

Sundararajan, Mukund: Talk Tue in Deep learning 7: analysis, Poster Tue #129

Suresh, Ananda: Talk Tue in Ranking and preferences, Talk Tue in Information theory, Poster Tue #112, Poster Tue #95

Suresh, Ananda: Talk Tue in Information theory, Poster Tue #128

Sussillo, David: Talk Tue in Recurrent neural networks 2, Poster Tue #63

Sutton, Charles: Talk Wed in Language 2, Poster Wed #97

Suzumura, Shinya: Talk Tue in Sparsity 2, Poster Tue #135

Sweeney, Latanya: Invited Talk Mon in Darling Harbour Theater

Świrszcz, Grzegorz: Talk Mon in Deep learning 1: backprop, Poster Mon #9

Synnaeve, Gabriel: Workshop Thu in C4.6

Szabo, Zoltan: Talk Wed in Kernel methods, Poster Wed #111

Szepesvari, Csaba: Talk Mon in Online learning 1, Poster Mon #30

Szlak, Liran: Talk Mon in Online learning 1, Poster Mon #14

Szorenyi, Balazs: Talk Mon in Online learning 1, Poster Mon #6

Taddy, Matt: Talk Wed in Causal Inference 2, Poster Wed #127

Tai, Cheng: Talk Tue in Continuous optimization 6, Poster Tue #71

Takac, Martin: Talk Tue in Continuous optimization 5, Poster Tue #48

Takeuchi, Ichiro: Talk Tue in Sparsity 2, Poster Tue #135

Taly, Ankur: Talk Tue in Deep learning 7: analysis, Poster Tue #129

Tamar, Aviv: Talk Mon in Reinforcement learning 2, Poster Mon #43

Tan, Zilong: Talk Mon in Matrix factorization 1, Poster Mon #55

Tan, Jie: Talk Wed in Large scale learning, Poster Wed #137

Tandon, Rashish: Talk Tue in Information theory, Poster Tue #136

Tang, Haoran: Talk Mon in Reinforcement learning 2, Poster Mon #52

Tang, Zhihao Gavin: Talk Mon in Semisupervised and curriculum learning, Poster Mon #93

Tang, Jian: Talk Mon in Deep learning theory 2, Poster Mon #83

Tang, Junqi: Talk Mon in Continuous optimization 2, Poster Mon #69

Tao, Dacheng: Talk Tue in Learning theory 2, Poster Tue #97, Talk Wed in Deep learning 8: hardware, Poster Wed #44

Tarlow, Daniel: Talk Tue in ML and programming, Poster Tue #134

Tay, Charlene: Talk Mon in Semisupervised and curriculum learning, Poster Mon #120

Taylor, Matthew: Talk Mon in Reinforcement learning 3, Poster Mon #97

Taylor, James: Talk Wed in Time series, Poster Wed #87

Taylor, Gavin: Talk Wed in Distributed optimization, Poster Wed #28

Tegmark, Max: Talk Tue in Recurrent neural networks 2, Poster Tue #47

Telgarsky, Matus: Talk Tue in Learning theory 1, Poster Wed #4

Theodorou, Evangelos: Talk Wed in Continuous control, Talk Wed in Time series, Poster Wed #38, Poster Wed #103

Thiran, Patrick: Talk Tue in Sparsity 2, Poster Tue #143

Thomas, Philip: Talk Tue in Reinforcement learning 5, Poster Tue #33

Tian, Yuandong: Talk Tue in Continuous optimization 6, Poster Tue #85

Titsias, Michalis: Talk Mon in Probabilistic learning 2, Poster Mon #80

Togelius, Julian: Workshop Thu in C4.6

Tokui, Seiya: Talk Mon in Deep learning 1: backprop, Poster Mon #33, Talk Tue in Information theory, Poster Tue #144

Tompson, Jonathan: Talk Wed in Applications, Poster Wed #140

Torr, Phil: Talk Wed in Deep reinforcement learning 1, Poster Wed #82

Torres Martins, Andre Filipe: Workshop Fri in C4.5

Tosatto, Samuele: Talk Mon in Reinforcement learning 1, Poster Mon #28

Tosh, Christopher: Talk Mon in Active learning, Poster Mon #145

Trabelsi, Chiheb: Talk Tue in Recurrent neural networks 2, Poster Tue #39

Tran, Dustin: Workshop Thu in Parkside 1

Tran, Bach: Talk Mon in Continuous optimization 2, Poster Mon #60

Tresp, Volker: Talk Tue in Recurrent neural networks 4, Poster Tue #122

Tripuraneni, Nilesh: Talk Tue in Probabilistic inference 3, Poster Tue #117, Talk Wed in Monte Carlo methods 2, Poster Wed #65

Trischler, Adam: Talk Mon in Deep learning 4: learning to learn, Poster Tue #21

Trivedi, Rakshit: Talk Tue in Networks and relational learning, Poster Tue #59

Trivedi, Rakshit: Talk Mon in Reinforcement learning 2, Poster Mon #79

Tsakiris, Manolis: Talk Tue in Clustering 1, Poster Tue #61

Tschiatschek, Sebastian: Talk Wed in Combinatorial optimization 1, Poster Wed #72

Tsuda, Koji: Talk Mon in Continuous optimization 4, Talk Tue in Sparsity 2, Poster Tue #26, Poster Tue #135

Tu, Stephen: Talk Mon in Continuous optimization 2, Poster Mon #51

Tunys, Tomas: Talk Mon in Online learning 1, Poster Mon #30

Turner, Richard: Talk Tue in Recurrent neural networks 3, Poster Tue #77, Talk Wed in Monte Carlo methods 2, Poster Wed #65

Tzamos, Christos: Talk Tue in Probabilistic inference 1, Poster Tue #50

Ubaru, Shashanka: Talk Wed in High dimensional estimation, Poster Wed #32

Udupa, Raghavendra: Talk Tue in Metric learning, Poster Wed #16

Uesato, Jonathan: Talk Tue in ML and programming, Poster Tue #118

Umez, Yuta: Talk Tue in Sparsity 2, Poster Tue #135

Umlauf, Jonas: Talk Wed in Continuous control, Poster Wed #47

Unser, Michael: Talk Tue in Sparsity 2, Poster Tue #143

Uria, Benigno: Talk Wed in Deep reinforcement learning 2, Poster Wed #125

Urschel, John: Talk Mon in Probabilistic learning 2, Poster Mon #62

Urtasun, Raquel: Tutorial Sun in Tutorials Session C, Talk Tue in Metric learning, Poster Wed #14, Workshop Thu in C4.10

Ushiku, Yoshitaka: Talk Mon in Transfer and multitask learning, Poster Mon #131

Usunier, Nicolas: Talk Tue in Deep learning 6, Poster Wed #9, Workshop Thu in C4.6

Valera, Isabel: Talk Mon in Probabilistic learning 1, Poster Mon #29

Valiant, Gregory: Talk Mon in Matrix factorization 3, Poster Tue #10

Valiant, Greg: Talk Tue in Learning theory 1, Poster Wed #2

Valko, Michal: Talk Mon in Online learning 2, Poster Mon #81, Talk Tue in Probabilistic inference 2, Poster Tue #80

van den Oord, Aaron: Talk Mon in Deep generative models 1, Poster Mon #10, Poster Mon #18, Talk Wed in Deep reinforcement learning 1, Poster Wed #64

van der Schaar, Mihaela: Talk Wed in Healthcare, Poster Wed #34

van Hasselt, Hado: Talk Wed in Deep reinforcement learning 1, Poster Wed #91

Van Roy, Benjamin: Talk Mon in Reinforcement learning 1, Poster Mon #36

Vanschoren, Joaquin: Workshop Thu in C4.9

Varma, Manik: Talk Mon in Supervised learning 1, Poster Mon #24, Talk Tue in Metric learning, Poster Wed #16, Workshop Thu in C4.7

Varoquaux, Gael: Talk Mon in Probabilistic learning 3, Poster Mon #98

Varshney, Pramod: Talk Tue in Continuous optimization 7, Poster Tue #131

Varshney, Kush: Workshop Thu in C4.8

Vassilvitskii, Sergei: Talk Tue in Clustering 1, Poster Tue #45

Vasudevan, Vijay: Talk Wed in Deep reinforcement learning 2, Poster Wed #134

Vaswani, Sharan: Talk Mon in Online learning 1, Poster Mon #22

Vayatis, Nicolas: Talk Mon in Continuous optimization 1, Poster Mon #19

Vedaldi, Andrea: Talk Mon in Deep learning 2: invariances, Poster Mon #49

Velingker, Ameya: Talk Mon in Supervised learning 2, Poster Mon #66

Venkataraman, Shivaram: Talk Mon in Continuous optimization 2, Poster Mon #51

Venkatesh, Svetha: Talk Wed in Gaussian processes, Poster Wed #142

Venkatraman, Arun: Talk Wed in Structured prediction, Poster Wed #121

Vetrov, Dmitry: Talk Wed in Deep learning 9: probabilistic, Poster Wed #79

Vezhnevets, Alexander: Talk Wed in Deep reinforcement learning 2, Poster Wed #107

Vian, John: Talk Tue in Game theory and multiagents, Poster Tue #140

Vidal, Rene: Talk Tue in Clustering 1, Poster Tue #61

Villacampa-Calvo, Carlos: Talk Wed in Gaussian processes, Poster Wed #117

Villegas, Ruben: Talk Tue in Recurrent neural networks 4, Poster Tue #106

Vinyals, Oriol: Tutorial Sun in Tutorials Session B, Talk Mon in Deep learning 1: backprop, Talk Mon in Deep generative models 1, Poster Mon #9, Poster Mon #1, Poster Mon #18, Talk Wed in Deep reinforcement learning 2, Talk Wed in Applications, Poster Wed #125, Poster Wed #131, Workshop Thu in C4.6

Vishwanath, Sriram: Talk Wed in Causal Inference 1, Poster Wed #84

Vladu, Adrian: Talk Mon in Continuous optimization 1, Poster Mon #3

Vollgraf, Roland: Talk Mon in Deep generative models 1, Poster Mon #34

Vorontsov, Eugene: Talk Tue in Recurrent neural networks 2, Poster Tue #39

Wahba, Grace: Talk Wed in Healthcare, Poster Wed #25

Wainwright, Martin: Talk Mon in Deep learning theory 3, Poster Mon #101

Walder, Christian: Talk Mon in Bayesian Nonparametrics, Poster Mon #141

Wang, Poan: Talk Mon in Matrix factorization 3, Poster Tue #15

Wang, Ziyu: Talk Mon in Deep generative models 1, Poster Mon #10

Wang, Shen: Talk Mon in Supervised learning 1, Poster Mon #39

Wang, Houfeng: Talk Mon in Deep learning 1: backprop, Poster Mon #17

Wang, Yixin: Talk Mon in Probabilistic learning 1, Poster Mon #21

Wang, Jialei: Talk Mon in Continuous optimization 3, Poster Mon #96, Talk Tue in Sparsity 2, Poster Tue #119

Wang, James: Talk Wed in Monte Carlo methods 1, Poster Wed #48

Wang, Jianmin: Talk Mon in Deep learning 3: metalearning, Poster Mon #85

Wang, Yanzhi: Talk Mon in Deep learning theory 2, Poster Mon #83

Wang, Yusu: Talk Mon in Probabilistic learning 3, Poster Mon #116

Wang, Chong: Talk Tue in Recurrent neural networks 4, Poster Tue #130

Wang, Yining: Talk Tue in Recurrent neural networks 4, Poster Tue #130, Talk Wed in Combinatorial optimization 2, Poster Wed #124

Wang, Zizhou: Talk Mon in Continuous optimization 1, Poster Mon #27

Wang, Yu-Xiang: Talk Tue in Reinforcement learning 5, Poster Tue #49

Wang, Taifeng: Talk Wed in Distributed optimization, Poster Wed #19

Wang, Guan: Talk Mon in Reinforcement learning 3, Poster Mon #97

Wang, Joseph: Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #54

Wang, Liwei: Talk Wed in Privacy and security 2, Poster Wed #102

Wang, Lezi: Talk Tue in Sparsity 1, Poster Tue #82

Wang, Yixin: Talk Mon in Probabilistic learning 1, Poster Mon #5

Wang, Yunhe: Talk Wed in Deep learning 8: hardware, Poster Wed #44

Wang, Mengdi: Talk Mon in Continuous optimization 1, Poster Mon #27

Wang, Di: Talk Tue in Spectral methods, Poster Tue #60

Wang, Shusen: Talk Tue in Learning theory 1, Poster Wed #1

Wang, Lingxiao: Talk Mon in Matrix factorization 3, Poster Tue #20, Talk Wed in Robust Estimation, Poster Wed #94, Poster Wed #76

Wang, Jie: Talk Tue in Sparsity 2, Poster Tue #111

Wang, Zi: Talk Mon in Bayesian Nonparametrics, Poster Mon #133, Talk Wed in Bayesian Optimization, Poster Wed #31

Wang, Yichen: Talk Tue in Networks and relational learning, Poster Tue #59, Talk Wed in Time series, Poster Wed #103

Watanabe, Shinji: Talk Wed in Language 1, Poster Wed #54

WEI, PENGFEI: Talk Mon in Transfer and multitask learning, Poster Mon #135

Weinberger, Kilian: Talk Tue in Deep learning 7: analysis, Poster Tue #137

Weiss, Ron: Talk Tue in Recurrent neural networks 3, Poster Tue #70

Weller, Adrian: Talk Tue in Probabilistic inference 3, Poster Tue #117, Workshop Thu in C4.8, Workshop Fri in C4.7

Welling, Max: Talk Wed in Deep learning 9: probabilistic, Poster Wed #70

Wen, Tsung-Hsien: Talk Wed in Language 2, Poster Wed #62, Workshop Thu in C4.11

Wen, Zheng: Talk Mon in Online learning 1, Poster Mon #22, Poster Mon #30

Weng, Paul: Talk Mon in Online learning 1, Poster Mon #6

White, Martha: Talk Mon in Reinforcement learning 4, Poster Tue #8, Talk Wed in Kernel methods, Poster Wed #129

Whiteson, Shimon: Talk Wed in Deep reinforcement learning 1, Poster Wed #82

Wichrowska, Olga: Talk Mon in Deep learning 4: learning to learn, Poster Tue #11

Wierstra, Daan: Talk Wed in Deep reinforcement learning 2, Poster Wed #125

Willett, Rebecca: Talk Wed in Robust Estimation, Poster Wed #85

Williams, Grady: Talk Wed in Time series, Poster Wed #103

Wilson, Ashia: Talk Mon in Continuous optimization 2, Poster Mon #51

Winner, Kevin: Talk Tue in Probabilistic inference 1, Poster Tue #42

Wolf, Lior: Talk Tue in ML and programming, Poster Tue #110

Wong, Eric: Talk Mon in Continuous optimization 4, Poster Tue #7

Wong, Sam: Talk Mon in Continuous optimization 1, Poster Mon #3

Woodruff, David: Talk Mon in Matrix factorization 3, Poster Tue #29

Wrigley, Andrew: Talk Tue in Probabilistic inference 1, Poster Tue #66

Wu, Steven: Talk Tue in Learning theory 1, Poster Wed #3

Wu, Chao-Yuan: Talk Mon in Continuous optimization 3, Poster Mon #105

Wu, Yuxin: Talk Tue in Networks and relational learning, Poster Tue #51

Wu, Xi-Zhu: Talk Wed in Structured prediction, Poster Wed #139

Xia, Yingce: Talk Mon in Supervised learning 2, Poster Mon #48

Xiao, Lin: Talk Mon in Continuous optimization 3, Poster Mon #96, Talk Tue in Reinforcement learning 5, Poster Tue #41

Xie, Pengtao: Talk Tue in Probabilistic inference 3, Poster Tue #133, Poster Tue #141

Xie, Qi: Talk Mon in Semisupervised and curriculum learning, Poster Mon #102

Xing, Eric: Talk Mon in Probabilistic learning 1, Poster Mon #13, Talk Tue in Probabilistic inference 3, Poster Tue #133, Poster Tue #141, Talk Wed in Language 2, Poster Wed #80

Xu, Chao: Talk Wed in Deep learning 8: hardware, Poster Wed #44

Xu, Chang: Talk Wed in Deep learning 8: hardware, Poster Wed #44

Xu, Yi: Talk Mon in Continuous optimization 1, Poster Mon #35

Xu, Huan: Talk Mon in Reinforcement learning 2, Poster Mon #79

Xu, Zheng: Talk Wed in Distributed optimization, Poster Wed #28

Xu, Lei: Talk Tue in Deep learning 6, Poster Wed #5

Xu, Pan: Talk Wed in Causal Inference 1, Poster Wed #66

Xu, Hongteng: Talk Wed in Time series, Poster Wed #78

Yadlowsky, Steven: Talk Tue in Continuous optimization 6, Poster Tue #99

Yan, Xinyan: Talk Wed in Continuous control, Poster Wed #38

Yan, Qiqi: Talk Tue in Deep learning 7: analysis, Poster Tue #129

Yang, Yinchong: Talk Tue in Recurrent neural networks 4, Poster Tue #122

Yang, Jimei: Talk Tue in Recurrent neural networks 4, Poster Tue #106

Yang, Bishan: Talk Wed in Structured prediction, Poster Wed #130

Yang, Bo: Talk Tue in Clustering 1, Poster Tue #53

Yang, Hongyu: Talk Mon in Probabilistic learning 2, Poster Mon #53

Yang, Yiming: Talk Tue in Networks and relational learning, Poster Tue #51

Yang, Jiachen: Talk Mon in Reinforcement learning 2, Poster Mon #79

Yang, Scott: Talk Mon in Deep learning 3: metalearning, Poster Mon #121, Workshop Fri in C4.1

Yang, Tianbao: Talk Mon in Continuous optimization 1, Talk Mon in Continuous optimization 3, Poster Mon #35, Poster Mon #114

Yang, Eunho: Talk Mon in Probabilistic learning 2, Poster Mon #44, Talk Tue in Sparsity 1, Poster Tue #96

Yang, Haichuan: Talk Tue in Sparsity 1, Poster Tue #89

Yang, Zichao: Talk Wed in Language 2, Talk Wed in Language 3, Poster Wed #80, Poster Wed #123

Yang, Zhuoran: Talk Wed in High dimensional estimation, Poster Wed #50

Yang, Lin: Talk Mon in Latent feature models, Poster Mon #31, Talk Tue in Clustering 2, Poster Wed #12

Yarats, Denis: Talk Wed in Language 3, Poster Wed #114

Yau, Christopher: Talk Mon in Probabilistic learning 2, Poster Mon #80

Ye, Nan: Talk Tue in Probabilistic inference 1, Poster Tue #66

Ye, Haishan: Talk Mon in Continuous optimization 4, Poster Tue #12

Ye, Yinyu: Talk Mon in Continuous optimization 1, Poster Mon #27

Ye, Xiaojing: Talk Mon in Reinforcement learning 2, Poster Mon #79

Ye, Jianbo: Talk Wed in Monte Carlo methods 1, Poster Wed #48

Ye, Jieping: Talk Tue in Sparsity 2, Poster Tue #111

Yen, En-Hsu: Talk Mon in Latent feature models, Talk Mon in Continuous optimization 3, Poster Mon #105, Poster Mon #23

Yin, Hao: Talk Mon in Continuous optimization 1, Poster Mon #27

yoon, jaehong: Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #62

Young, Stephen: Talk Wed in Language 2, Poster Wed #62

Yu, Yaoliang: Talk Tue in Probabilistic inference 3, Poster Tue #141

Yu, Nenghai: Talk Mon in Supervised learning 2, Poster Mon #48, Talk Wed in Distributed optimization, Poster Wed #19

Yu, Qi: Workshop Fri in C4.1

Yu, Felix: Talk Tue in Information theory, Poster Tue #112

Yu, Hong: Talk Mon in Deep learning 3: metalearning, Poster Mon #94

Yu, Chen: Talk Mon in Semisupervised and curriculum learning, Poster Mon #120

Yu, Philip: Talk Mon in Supervised learning 1, Poster Mon #39

Yuan, Xiaoming: Talk Wed in Distributed optimization, Poster Wed #28

Yuan, Xiaotong: Talk Tue in Sparsity 1, Poster Tue #82

Yuan, Bo: Talk Mon in Deep learning theory 2, Poster Mon #83

Yuan, Ganzhao: Talk Mon in Continuous optimization 2, Poster Mon #42

Yue, Yisong: Talk Tue in Game theory and multiagents, Poster Tue #132

Yue, Shichao: Talk Wed in Healthcare, Poster Wed #60

Yurochkin, Mikhail: Talk Tue in Clustering 1, Poster Tue #69

Zadimoghaddam, Morteza: Talk Wed in Combinatorial optimization 1, Poster Wed #90

Zahavy, Tom: Workshop Thu in C4.1

Zaheer, Manzil: Talk Tue in Recurrent neural networks 4, Poster Tue #138, Talk Wed in Monte Carlo methods 1, Poster Wed #39

Zandieh, Amir: Talk Mon in Supervised learning 2, Poster Mon #66

Zappella, Giovanni: Talk Mon in Online learning 3, Poster Mon #99

Zenke, Friedemann: Talk Tue in Deep learning 5: Fisher approximations, Poster Tue #46

Zha, Hongyuan: Talk Mon in Reinforcement learning 2, Poster Mon #79, Talk Wed in Time series, Poster Wed #78

Zhai, Chengxiang: Talk Wed in Robust Estimation, Poster Wed #94

Zhang, Chicheng: Workshop Fri in C4.6

Zhang, Zhihua: Talk Mon in Continuous optimization 4, Poster Tue #12

Zhang, Chicheng: Talk Mon in Online learning 4, Poster Tue #9

Zhang, Chenzi: Talk Mon in Semisupervised and curriculum learning, Poster Mon #93

Zhang, Hongyang: Talk Wed in Privacy and security 2, Poster Wed #86

Zhang, Yuchen: Talk Mon in Deep learning theory 3, Poster Mon #101

Zhang, Lijun: Talk Mon in Continuous optimization 3, Poster Mon #114

Zhang, Hantian: Talk Wed in Large scale learning, Poster Wed #128

Zhang, Weizhong: Talk Tue in Sparsity 2, Poster Tue #111

Zhang, Ce: Talk Wed in Large scale learning, Poster Wed #128

Zhang, Wenpeng: Talk Wed in Distributed optimization, Poster Wed #46

Zhang, Huan: Talk Tue in Ensemble methods, Poster Tue #79

Zhang, Marvin: Talk Wed in Continuous control, Poster Wed #29

Zhang, Yi: Talk Mon in Deep generative models 2, Poster Mon #41

Zhang, Xiao: Talk Wed in Gaussian processes, Poster Wed #135

Zhang, Cyril: Talk Mon in Online learning 2, Poster Mon #45

Zhang, Tong: Talk Tue in Sparsity 2, Poster Tue #119, Talk Wed in Distributed optimization, Poster Wed #46

Zhang, Cheng: Talk Wed in Monte Carlo methods 2, Poster Wed #74

Zhang, Xiao: Talk Mon in Matrix factorization 3, Poster Tue #20

Zhang, Shengyu: Talk Tue in Probabilistic inference 3, Poster Tue #125

Zhang, Yizhe: Talk Wed in Monte Carlo methods 2, Talk Wed in Language 2, Poster Wed #83, Poster Wed #89

Zhang, Qin: Talk Mon in Online learning 3, Poster Mon #125

Zhang, Ningshan: Workshop Fri in C4.6

Zhang, Yilin: Talk Wed in Healthcare, Poster Wed #25

Zhang, Teng: Talk Mon in Supervised learning 1, Poster Mon #32

Zhao, Mingmin: Talk Wed in Healthcare, Poster Wed #60

Zhao, Tuo: Talk Mon in Latent feature models, Poster Mon #31

Zhao, Liang: Talk Mon in Deep learning theory 2, Poster Mon #83

Zhao, Ming-Min: Talk Mon in Continuous optimization 3, Poster Mon #87

Zhao, Shengjia: Talk Mon in Deep generative models 3, Poster Mon #86

Zhao, He: Talk Tue in Networks and relational learning, Poster Tue #35

Zhao, Peilin: Talk Wed in Distributed optimization, Poster Wed #46

Zhe, Shandian: Talk Wed in Gaussian processes, Poster Wed #135

Zheng, Shuxin: Talk Wed in Distributed optimization, Poster Wed #19

Zheng, Kai: Talk Wed in Privacy and security 2, Poster Wed #102

Zheng, Shuai: Talk Mon in Deep learning theory 2, Poster Mon #74

Zheng, Yi: Talk Mon in Matrix factorization 2, Poster Mon #109

Zhong, Kai: Talk Mon in Deep learning theory 2, Poster Mon #47

Zhou, Yuefeng: Talk Wed in Deep learning 8: hardware, Poster Wed #17

Zhou, Dengyong: Talk Mon in Online learning 3, Poster Mon #108, Talk Tue in Reinforcement learning 5, Talk Tue in Recurrent neural networks 4, Poster Tue #130, Poster Tue #41

Zhou, Zhi-Hua: Talk Mon in Supervised learning 1, Poster Mon #32

Zhou, Mingyuan: Talk Wed in Monte Carlo methods 2, Poster Wed #100

Zhou, Yi: Talk Tue in Continuous optimization 5, Poster Tue #40

Zhou, Yichi: Talk Mon in Online learning 2, Poster Mon #54

Zhou, Yuan: Talk Mon in Online learning 3, Poster Mon #125

Zhou, Chaoxu: Talk Mon in Continuous optimization 4, Poster Tue #17

Zhou, Yi: Talk Tue in Continuous optimization 7, Talk Tue in Probabilistic inference 3, Poster Tue #141, Poster Tue #131

Zhou, Hao: Talk Wed in Healthcare, Poster Wed #25

Zhou, Zhi-Hua: Talk Wed in Structured prediction, Poster Wed #139

Zhu, Han: Talk Mon in Deep learning 3: metalearning, Poster Mon #85

Zhu, Zhenyao: Talk Wed in Language 1, Poster Wed #27

Zhu, Rongda: Talk Wed in Robust Estimation, Poster Wed #94

Zhu, Junyan: Workshop Thu in C4.3

zhu, wenwu: Talk Wed in Distributed optimization, Poster Wed #46

Zhu, Jun: Talk Mon in Online learning 2, Poster Mon #54

Zilly, Julian: Talk Mon in Recurrent neural networks 1, Poster Mon #140

Zink, Daniel: Talk Tue in Continuous optimization 5, Poster Tue #32, Poster Tue #40

Zoghi, Masrour: Talk Mon in Online learning 1, Poster Mon #30

Zoph, Barret: Talk Wed in Deep reinforcement learning 2, Poster Wed #134

Zou, Yuliang: Talk Tue in Recurrent neural networks 4, Poster Tue #106

Zou, James: Talk Tue in Learning theory 1, Talk Tue in Probabilistic inference 3, Poster Tue #141, Poster Wed #2